



US007668095B2

(12) **United States Patent**
Chase et al.

(10) **Patent No.:** **US 7,668,095 B2**

(45) **Date of Patent:** ***Feb. 23, 2010**

(54) **TRAFFIC MANAGEMENT FOR FRAME RELAY SWITCHED DATA SERVICE**

(75) Inventors: **Christopher J. Chase**, Freehold, NJ (US); **Stephen L. Holmgren**, Little Silver, NJ (US); **John Babu Medamana**, Colts Neck, NJ (US); **Vikram R. Saksena**, Freehold, NJ (US)

(73) Assignee: **AT&T Corp.**, New York, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 696 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **11/019,330**

(22) Filed: **Dec. 21, 2004**
(Under 37 CFR 1.47)

(65) **Prior Publication Data**
US 2005/0105466 A1 May 19, 2005

Related U.S. Application Data

(63) Continuation of application No. 09/654,616, filed on Sep. 1, 2000, now Pat. No. 6,847,611, which is a continuation of application No. 08/988,424, filed on Dec. 10, 1997, now Pat. No. 6,188,671.

(60) Provisional application No. 60/051,564, filed on Jul. 3, 1997.

(51) **Int. Cl.**
H04L 1/00 (2006.01)

(52) **U.S. Cl.** **370/232; 370/229**

(58) **Field of Classification Search** **370/231, 370/232, 233, 234, 235, 236, 235.1, 236.2, 370/252, 253, 258, 360, 386, 389, 391, 396, 370/395.21, 395.3, 395.41, 395.43, 395.64, 370/229**

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,018,993 A 4/1977 Edstrom

(Continued)

FOREIGN PATENT DOCUMENTS

CA 2202542 10/1997

(Continued)

OTHER PUBLICATIONS

Fred Halsall, Data Communications, Computer Networks and Open Systems, Fourth Ed. (England: Addison-Wesley Publishers Ltd. 1996), 470-475.

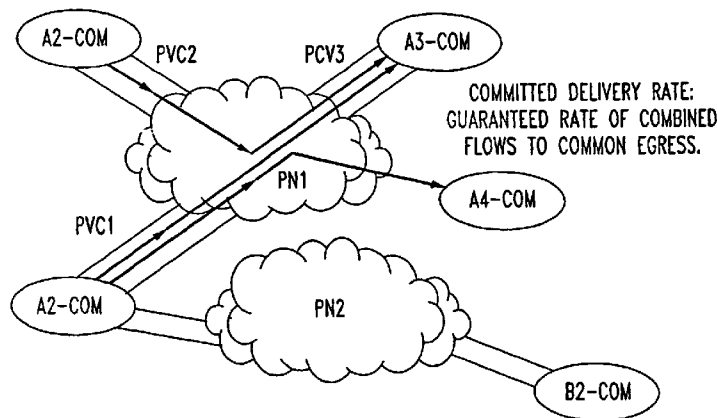
(Continued)

Primary Examiner—Brenda Pham

(57) **ABSTRACT**

A new type of data transport service which uses a frame relay layer 2 data link connection identifier (DLCI) to select among various service types, feature sets, and/or closed user groups (CUGs). A layer 3 address may be extracted from a layer 2 frame, and the layer 3 address information may be used to route a data packet over a packet-switched network according to the service classes, feature sets, and/or CUGs selected. At the destination, the layer 3 data packet may again be enclosed in a layer 2 frame with a DLCI indicating the service classes, features sets, and/or CUGs. Because the use of conventional permanent virtual circuits (PVCs) is not required in aspects of the invention, new methods of measuring and managing network traffic are presented.

8 Claims, 10 Drawing Sheets



U.S. PATENT DOCUMENTS						
			5,058,138	A	10/1991	Figura et al.
			5,065,397	A	11/1991	Shlobara
4,135,156	A	1/1979	5,079,764	A	1/1992	Orita et al.
4,181,886	A	1/1980	5,081,621	A	1/1992	Sugimoto
4,285,064	A	8/1981	5,086,426	A	2/1992	Tsukakoshi et al.
4,301,533	A	11/1981	5,090,011	A	2/1992	Fukuta et al.
4,307,461	A	12/1981	5,095,480	A	3/1992	Fenner
4,319,353	A	3/1982	5,101,267	A	3/1992	Morales-Garza
4,320,504	A	3/1982	5,115,433	A	5/1992	Baran et al.
4,328,543	A	5/1982	5,121,396	A	6/1992	Irvin et al.
4,330,857	A	5/1982	5,130,978	A	7/1992	Mobasser
4,332,026	A	5/1982	5,153,876	A	10/1992	Sin
4,346,470	A	8/1982	5,177,604	A	1/1993	Martinez
4,377,793	A	3/1983	5,184,347	A	2/1993	Farwell et al.
4,381,562	A	4/1983	5,191,582	A	3/1993	Upp
4,468,727	A	8/1984	5,195,090	A	3/1993	Bolliger et al.
4,485,478	A	11/1984	5,195,091	A	3/1993	Farwell et al.
4,507,781	A	3/1985	5,208,811	A	5/1993	Kashio et al.
4,516,156	A	5/1985	5,223,923	A	6/1993	Morales-Garza
4,521,879	A	6/1985	5,224,099	A	6/1993	Corbalis et al.
4,536,874	A	8/1985	5,229,990	A	7/1993	Teraslinna
4,587,651	A	5/1986	5,229,994	A	7/1993	Balzano et al.
4,597,077	A	6/1986	5,237,564	A	8/1993	Lespagnol et al.
4,599,647	A	7/1986	5,239,545	A	8/1993	Buchholz
4,642,806	A	2/1987	5,251,207	A	10/1993	Abensour et al.
4,644,534	A	2/1987	5,260,936	A	11/1993	Bardet et al.
4,665,516	A	5/1987	5,274,643	A	12/1993	Fisk
4,672,602	A	6/1987	5,276,907	A	1/1994	Meldan
4,679,191	A	7/1987	5,278,892	A	1/1994	Bollinger et al.
4,686,698	A	8/1987	5,280,481	A	1/1994	Chang et al.
4,701,907	A	10/1987	5,287,535	A	2/1994	Sakagawa et al.
4,703,479	A	10/1987	5,291,479	A	3/1994	Vaziri et al.
4,706,080	A	11/1987	5,291,492	A	3/1994	Andrews et al.
4,706,081	A	11/1987	5,295,133	A	3/1994	Jurkevich
4,710,917	A	12/1987	5,295,137	A	3/1994	Jurkevich
4,713,806	A	12/1987	5,301,273	A	4/1994	Konishi
4,720,850	A	1/1988	5,303,042	A	4/1994	Lewis et al.
4,720,873	A	1/1988	5,303,237	A	4/1994	Bergman et al.
4,730,305	A	3/1988	5,305,308	A	4/1994	English et al.
4,739,510	A	4/1988	5,307,348	A	4/1994	Buchholz et al.
4,751,732	A	6/1988	5,309,440	A	5/1994	Nakamura et al.
4,769,833	A	9/1988	5,313,454	A	5/1994	Daley et al.
4,775,974	A	10/1988	5,313,457	A	5/1994	Hostetter et al.
4,777,657	A	10/1988	5,315,591	A	5/1994	Brent et al.
4,792,948	A	12/1988	5,319,705	A	6/1994	Halter et al.
4,797,589	A	1/1989	5,319,707	A	6/1994	Wasilewski et al.
4,797,878	A	1/1989	5,325,421	A	6/1994	Hou et al.
4,811,338	A	3/1989	5,327,421	A	7/1994	Hiller et al.
4,813,039	A	3/1989	5,331,631	A	7/1994	Teraslinna
4,847,829	A	7/1989	5,335,276	A	8/1994	Thompson et al.
4,847,892	A	7/1989	5,337,044	A	8/1994	Folger et al.
4,858,225	A	8/1989	5,343,465	A *	8/1994	Khalil 370/232
4,868,811	A	9/1989	5,345,445	A	9/1994	Hiller et al.
4,876,737	A	10/1989	5,345,446	A	9/1994	Hiller et al.
4,879,711	A	11/1989	5,351,146	A	9/1994	Chan et al.
4,888,769	A	12/1989	5,355,362	A	10/1994	Gorshe et al.
4,890,280	A	12/1989	5,361,259	A	11/1994	Hunt et al.
4,894,822	A	1/1990	5,363,369	A	11/1994	Hemmady et al.
4,916,691	A	4/1990	5,365,265	A	11/1994	Shibata et al.
4,933,936	A	6/1990	5,365,524	A	11/1994	Hiller et al.
4,937,825	A	6/1990	5,373,550	A	12/1994	Campbell et al.
RE33,426	E	11/1990	5,381,408	A	1/1995	Brent et al.
4,970,721	A	11/1990	5,387,941	A	2/1995	Montgomery et al.
4,993,015	A	2/1991	5,390,169	A	2/1995	Bernstein
4,999,829	A	3/1991	5,390,175	A	2/1995	Hiller et al.
5,014,267	A	5/1991	5,392,277	A	2/1995	Bernstein
5,016,243	A	5/1991	5,392,402	A	2/1995	Robrock, II
5,018,133	A	5/1991	5,394,561	A	2/1995	Freeburg
5,019,910	A	5/1991	5,400,329	A	3/1995	Tokura et al.
5,023,869	A	6/1991	5,400,401	A	3/1995	Wasilewski et al.
5,023,873	A	6/1991	5,402,418	A	3/1995	Shibata et al.
5,027,400	A	6/1991	5,402,478	A	3/1995	Hluchyj et al.
5,029,163	A	7/1991	5,404,505	A	4/1995	Levinson

5,408,237 A	4/1995	Patterson et al.	5,768,271 A	6/1998	Seid et al.
5,408,419 A	4/1995	Wong et al.	5,796,727 A	8/1998	Harrison et al.
5,408,464 A	4/1995	Jurkevich	5,796,944 A	8/1998	Hill et al.
5,412,647 A	5/1995	Giroux et al.	5,802,365 A	9/1998	Kathail et al.
5,412,660 A	5/1995	Chen et al.	5,805,599 A	9/1998	Mishra et al.
5,416,770 A	5/1995	Stoner et al.	5,809,025 A	9/1998	Timbs
5,420,857 A	5/1995	Jurkevich	5,812,527 A	9/1998	Kline et al.
5,420,866 A	5/1995	Wasllewski	5,812,545 A	9/1998	Liebowitz et al.
5,422,882 A	6/1995	Hiller et al.	5,815,495 A	9/1998	Saitoh et al.
5,423,002 A	6/1995	Hart	5,828,666 A	10/1998	Focsaneanu et al.
5,425,029 A	6/1995	Hluchyj et al.	5,852,721 A	12/1998	Dillon et al.
5,426,635 A	6/1995	Mitra et al.	5,878,045 A	3/1999	Timbs
5,426,636 A	6/1995	Hiller et al.	5,883,893 A	3/1999	Rumer
5,426,640 A	6/1995	Hluchyj et al.	5,903,559 A	5/1999	Acharya et al.
5,426,773 A	6/1995	Chabanet et al.	5,907,677 A	5/1999	Glenn et al.
5,428,607 A	6/1995	Hiller et al.	5,909,443 A *	6/1999	Fichou et al. 370/412
5,428,608 A	6/1995	Freeman et al.	5,914,956 A	6/1999	Williams
5,428,616 A	6/1995	Field et al.	5,920,560 A	7/1999	Takahashi et al.
5,428,629 A	6/1995	Gutman et al.	5,930,259 A	7/1999	Katsube et al.
5,430,727 A	7/1995	Callon	5,946,313 A	8/1999	Allan et al.
5,434,850 A	7/1995	Fielding et al.	5,970,048 A	10/1999	Pajuvirta et al.
5,434,856 A	7/1995	Huang	5,974,052 A	10/1999	Johnson et al.
5,436,891 A	7/1995	Grossman et al.	5,982,783 A	11/1999	Frey et al.
5,438,565 A	8/1995	Hemmady et al.	5,982,869 A	11/1999	Lozano
5,442,691 A	8/1995	Price et al.	5,991,268 A	11/1999	Awdeh et al.
5,446,726 A	8/1995	Rostoker et al.	5,995,511 A	11/1999	Zhou et al.
5,446,730 A	8/1995	Lee et al.	5,999,610 A	12/1999	Lin et al.
5,446,733 A *	8/1995	Tsuruoka 370/232	6,009,097 A	12/1999	Han
5,448,564 A	9/1995	Thor	6,011,778 A	1/2000	Kilkki et al.
5,450,394 A	9/1995	Gruber et al.	6,012,090 A	1/2000	Chung
5,450,395 A	9/1995	Hostetter et al.	6,018,530 A	1/2000	Chakrovorty
5,452,293 A	9/1995	Wilkinson et al.	6,023,453 A	2/2000	Ruutu et al.
5,453,979 A	9/1995	Schibler et al.	6,023,470 A	2/2000	Lee et al.
5,457,808 A	10/1995	Osawa et al.	6,081,524 A	6/2000	Chase et al.
5,457,811 A	10/1995	Lemson	6,108,336 A	8/2000	Duault
5,459,723 A	10/1995	Thor	6,115,748 A	9/2000	Hauser
5,463,620 A	10/1995	Sriram	6,118,759 A	9/2000	Ohyoshi et al.
5,463,777 A	10/1995	Bialkowski et al.	6,151,324 A	11/2000	Belser et al.
5,469,468 A	11/1995	Schilling	6,185,215 B1	2/2001	Aho
5,469,495 A	11/1995	Beveridge	6,188,671 B1 *	2/2001	Chase et al. 370/232
5,471,473 A	11/1995	Tejima	6,212,163 B1	4/2001	Aida
5,473,612 A	12/1995	Dehner, Jr. et al.	6,230,203 B1	5/2001	Koperda et al.
5,481,561 A	1/1996	Fang	6,278,693 B1	8/2001	Aldred et al.
5,481,609 A	1/1996	Cohen et al.	6,298,120 B1	10/2001	Civanlar et al.
5,490,140 A	2/1996	Abensour	6,359,909 B1 *	3/2002	Ito et al. 370/522
5,490,252 A	2/1996	Macera et al.	6,400,681 B1	6/2002	Bertin et al.
5,491,695 A	2/1996	Meagher et al.	6,487,217 B1	11/2002	Baroudi
5,495,607 A	2/1996	Pisello et al.	6,526,023 B1	2/2003	Koga et al.
5,497,375 A	3/1996	Hluchyj	6,847,611 B1 *	1/2005	Chase et al. 370/232
5,506,904 A	4/1996	Sheldrick et al.	6,870,827 B1	3/2005	Voit et al.
5,517,502 A	5/1996	Bestler et al.	6,904,037 B2	6/2005	Oran et al.
5,519,699 A	5/1996	Ohsawa	6,967,922 B1	11/2005	Ohyoshi et al.
5,521,914 A	5/1996	Mavraganis et al.	2003/0161328 A1	8/2003	Chase et al.
5,526,404 A	6/1996	Wiedeman et al.	2006/0104273 A1	5/2006	Chase et al.
5,526,489 A	6/1996	Nilakantan et al.	2007/0253415 A1	11/2007	Chase et al.
5,539,884 A	7/1996	Robrock, II	2009/0041022 A1	2/2009	Chase et al.
5,541,922 A	7/1996	Pyhalammi			
5,565,910 A	10/1996	Rowse et al.			
5,566,173 A	10/1996	Steinbrecher			
5,579,370 A	11/1996	Fukushima et al.	EP	0 324 363	7/1989
5,587,743 A	12/1996	Montgomery et al.	EP	0 535 890	4/1993
5,590,133 A	12/1996	Billstrom et al.	EP	0 719 012 A2	6/1996
5,594,782 A	1/1997	Zicker et al.	EP	0 800 294 A1	10/1997
5,623,492 A *	4/1997	Teraslinna 370/397	WO	WO 96 19060 A	6/1996
5,628,011 A	5/1997	Ahamed et al.			
5,654,957 A	8/1997	Koyama			
5,675,576 A	10/1997	Kalampoukas			
5,675,577 A	10/1997	Komatsu			
5,689,553 A	11/1997	Ahuja et al.			
5,727,065 A	3/1998	Dillon			
5,734,654 A	3/1998	Shirai et al.			
5,743,530 A	4/1998	Sklad et al.			
5,754,530 A	5/1998	Awdeh et al.			

FOREIGN PATENT DOCUMENTS

EP	0 324 363	7/1989
EP	0 535 890	4/1993
EP	0 719 012 A2	6/1996
EP	0 800 294 A1	10/1997
WO	WO 96 19060 A	6/1996

OTHER PUBLICATIONS

Bonomi, F. et al., "The Rate-Base Flow Control Framework for the Available Bit Rate ATM Service", IEEE Network, Mar./Apr. 1995, pp. 25-39, vol. 9, No. 2, ISSN:0890-8044, IEEE Inc., New York, US.

Sharon, C.M. et al., "Modeling and Control of VBR H.261 Video Transmission over Frame Relay Networks", IEEE Transactions on Circuits and Systems for Video Technology, Jun. 1997, pp. 527-538, vol. 7, No. 3, ISSN:1051-8215, IEEE Inc.

- Beraldi, R. et al., "Selective BECN Schemes for Congestion Control of ABR Traffic in ATM LAN", 1996 IEEE International Conference on Communications (ICC), Converging Technologies for Tomorrow's Applications. Dallas, Jun. 23-27 1996, pp. 503-507, vol. 1 ISSN: 0-7803-3250-4, IEEE Inc. New York, US.
- EPO Search report, mail date Jul. 29, 2003, pp. 1-6 for EP 98123041.
- Bih-Hwang Lee et al., "A Performance Study on Frame Relay and ATM Interworking Unit and its Applications: Connectionless Data Service and VPN", IEEE Trans. Commun., vol. E80-B, No. 6, Jun. 1997, pp. 948-962.
- Peter Newman et al., "IP Switching and Gigabit Routers", IEEE Communications Magazine, Jan. 1997, pp. 64-69.
- Stephen Saunders, "Concentrated WAN Access to ATM and Frame Relay", Data Communications, vol. 25, No. 10 Aug. 1996, 2 pages.
- Klessig, Robert W. et al., "SMDS Wide-area Data Networking With Switched Multit-megabit Data Service", Prentice Hall, 1995 pp. 11-20.
- EPO Search Report, Nov. 20, 1998, 3 pages.
- Department of Computer Sciences, Purdue University, Douglas E. Comer, "Interworking with TCP/IP" vol. I: Principles, Protocols, and Architecture, Second Edition, 1991, 4 pages.
- Internet article, "Short-cut Routing", Feb. 14, 1997.
- "Active Networks New Networking System Slashes Costs and Unleashes Productivity!", True Networking Applications from Advanced Compression Technology, Inc. vol. 1, No. 1.
- "Reasonable Connectivity Through Multimedia Networking System", Technological Applications, pp. 30-32, Sep. 1992.
- "Multiplex Transmission in Carrier Telephony", Reference Data for Radio Engineers, Fifth Ed. pp. 30-23 through 30-27, 1968.
- "Contact—The SDM-HDLC Voice Packetizer Transmits Voice Over a Packet Network", Advanced Compression Technology, Inc., pp. 1-2.
- "Fast-module T-1/E-1 Performance for CPX Switch", Dynatech Communications, pp. 1-8.
- Briley, J.B., "MS-1000/MS-2000 Multimedia Switches", Certified Network Systems and Services, 1991.
- Cooper, N., "Network Technology for Multimedia", pp. 6/1-6/3s, Sep. 1992.
- Friend, G.E., et al., "Fiber Optic and Satellite Communications", Understanding Data Communications, pp. 1-168, 1998.
- Friend, G.E., et al., "Multiplexing Techniques", Understanding Data Communications, pp. 161-176, 1998.
- Friend, G.E., et al., "Fiber Optic and Satellite Communications", pp. 202-207, 1998.
- Friend, G.E., et al., "Satellite Communications Systems", Understanding Data Communications, pp. 243-270, 1998.
- Pinsky, D., "Total Compression—Frame Relay Net Runs Voice, Data", World News, 1993.
- Platt, A., et al., "Some Aspects of Traffic Management in Frame Relay Networks", pp. 23/1-23/6, Apr. 1991.
- "Model 710 Communications Subsystem", Harris Corporation, 1986.
- "Model 712 Communications Subsystem", Harris Corporation 1986.
- Collier, A., "Wellfleet Unveils Two Bridge-Routers Aimed at Branch Networking", Communications Week, 1993.
- Fleming, S., "What Users Can Expect from New Virtual Wideband Services", Wdwband Services, 1990, pp. 29-35.
- Gordon, G.D., et al., Modulation, Multiplexing, and Multiple Access, Principles of Communications Satellites, , 1993 pp. 114-135.
- Llana, Jr., A., "Bringing the Telephone to Mexico—Satellites Work in Rural Environment", Via Satellite, 1992, pp. 52-54.
- Malone, R., "Frame Relay Service Pricing: Contrasting the Competitors", Business Communications Review, 1993, pp. 32-37.
- Muller, N. J., "The Management Aspects of Frame Relay", International Journal of Network Management, vol. 2, 1992, pp. 87-99.
- Pontano, B.A., et al., "ISDN Over Satellite Networks", 1992, pp. 0560-0564.
- Rana, A.H., et al., "VSAT-Enhanced ISDN: Architectures and Implementation", IEEE Journal on Selected Area of Communications, 10(1992) August, No. 6.
- Roy, A.K., "Frame Relay and its Transportation Dilemma", Proc. 1991 Singaport Int. Conf. on Networks, 1991 pp. 393-400.
- Salamoff, S.B., et al., VSAT System Considerations for the ORION Satellites, 1992, pp. 1-7.
- Shum, M., Special Report: Satellite Transmission, SCPC: Satellite Links with Infinite Uses, Telecom Asia, 1993. pp. 9-10.
- * cited by examiner

FIG. 1

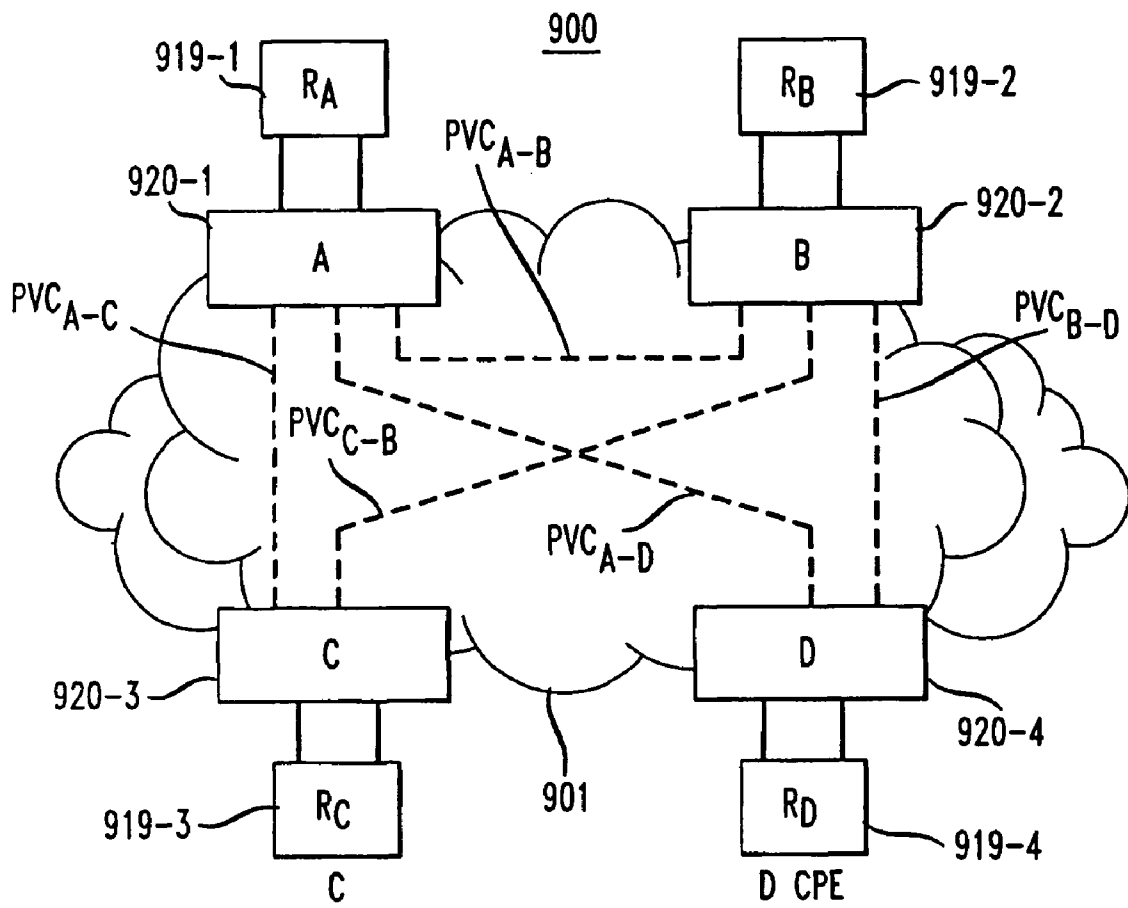


FIG. 2

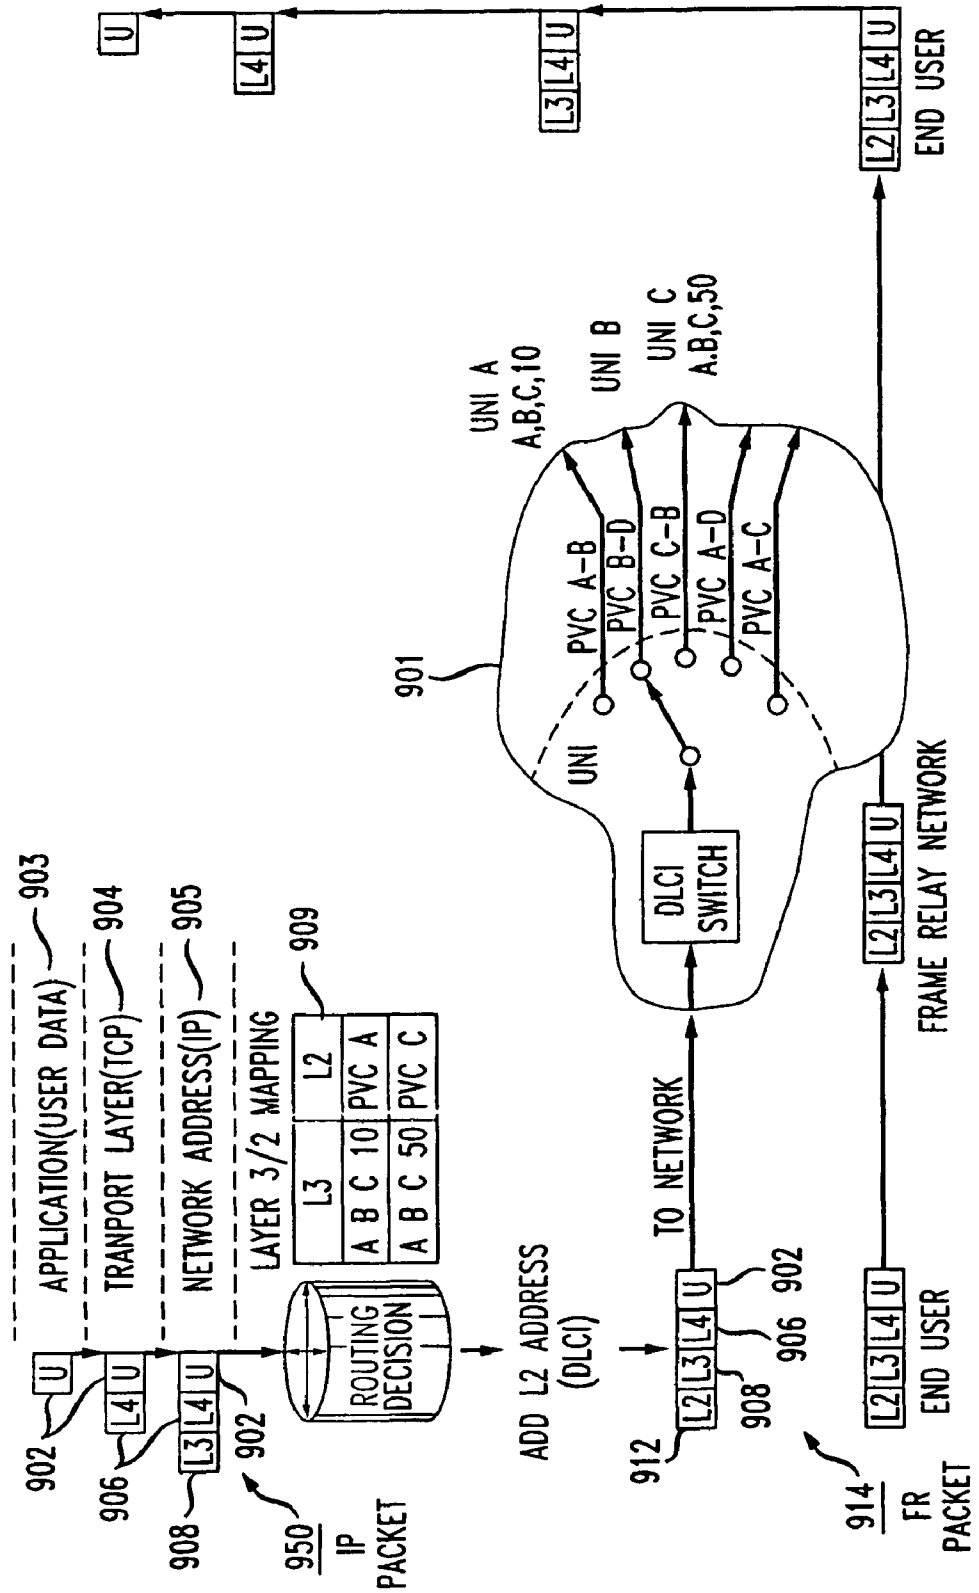


FIG. 3

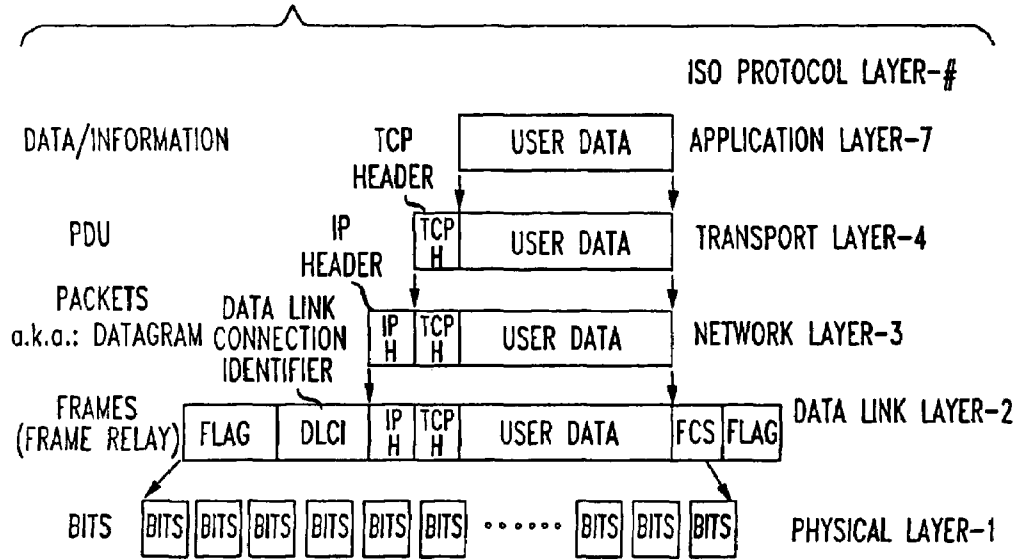
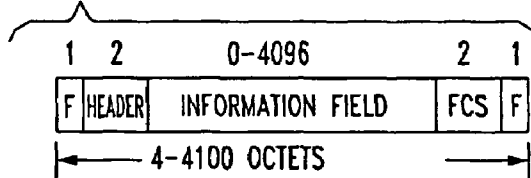
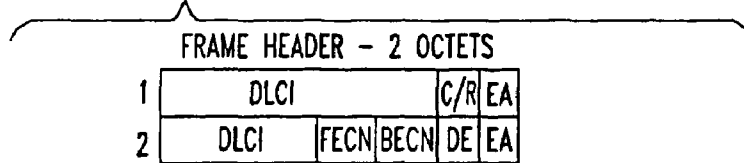


FIG. 4



F = FRAME DELIMITER (FLAG)
 (01111110 BINARY VALUE)
 HEADER = FR HEADER FIELD
 INFORMATION FIELD = USER'S PAYLOAD
 FCS = FRAME CHECK SEQUENCE

FIG. 5



DLCI = DATA LINK CONNECTION IDENTIFIER (10 BITS)
 C/R = COMMAND / RESPONSE (1BIT)
 EA = EXTENDED ADDRESS (2 BITS)
 FECN = FORWARD EXPLICIT CONGESTION NOTIFICATION (1 BIT)
 BECN = BACKWARD EXPLICIT CONGESTION NOTIFICATION (1 BIT)
 DE = DISCARD ELIGIBILITY (1 BIT)

FIG. 6

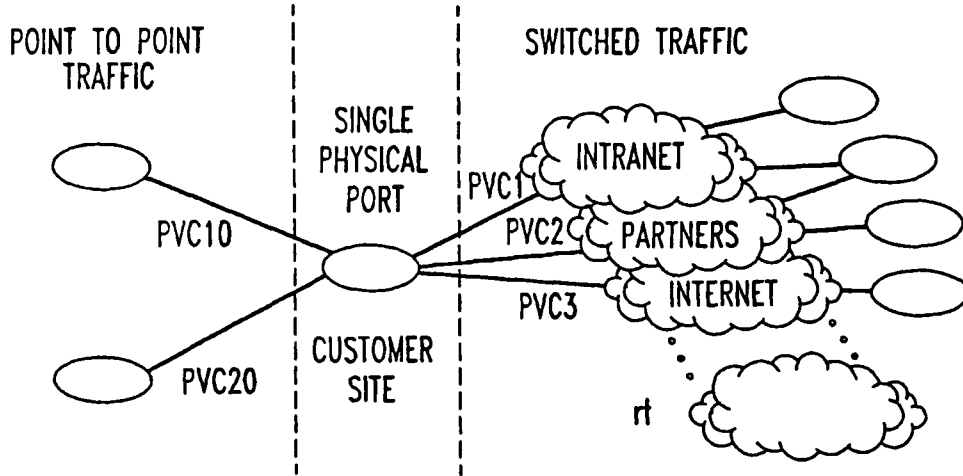


FIG. 11

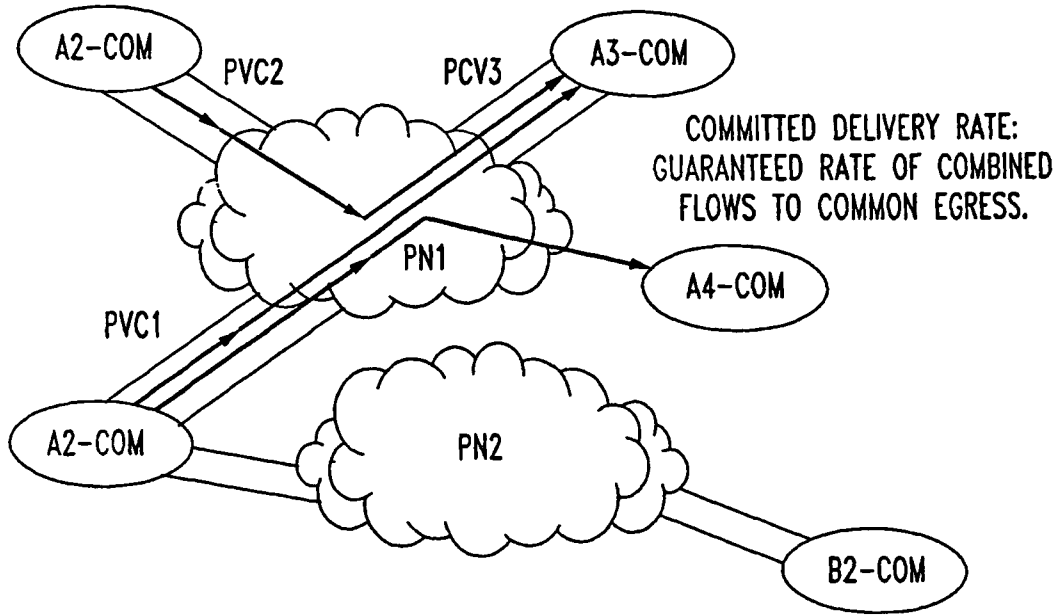


FIG. 7

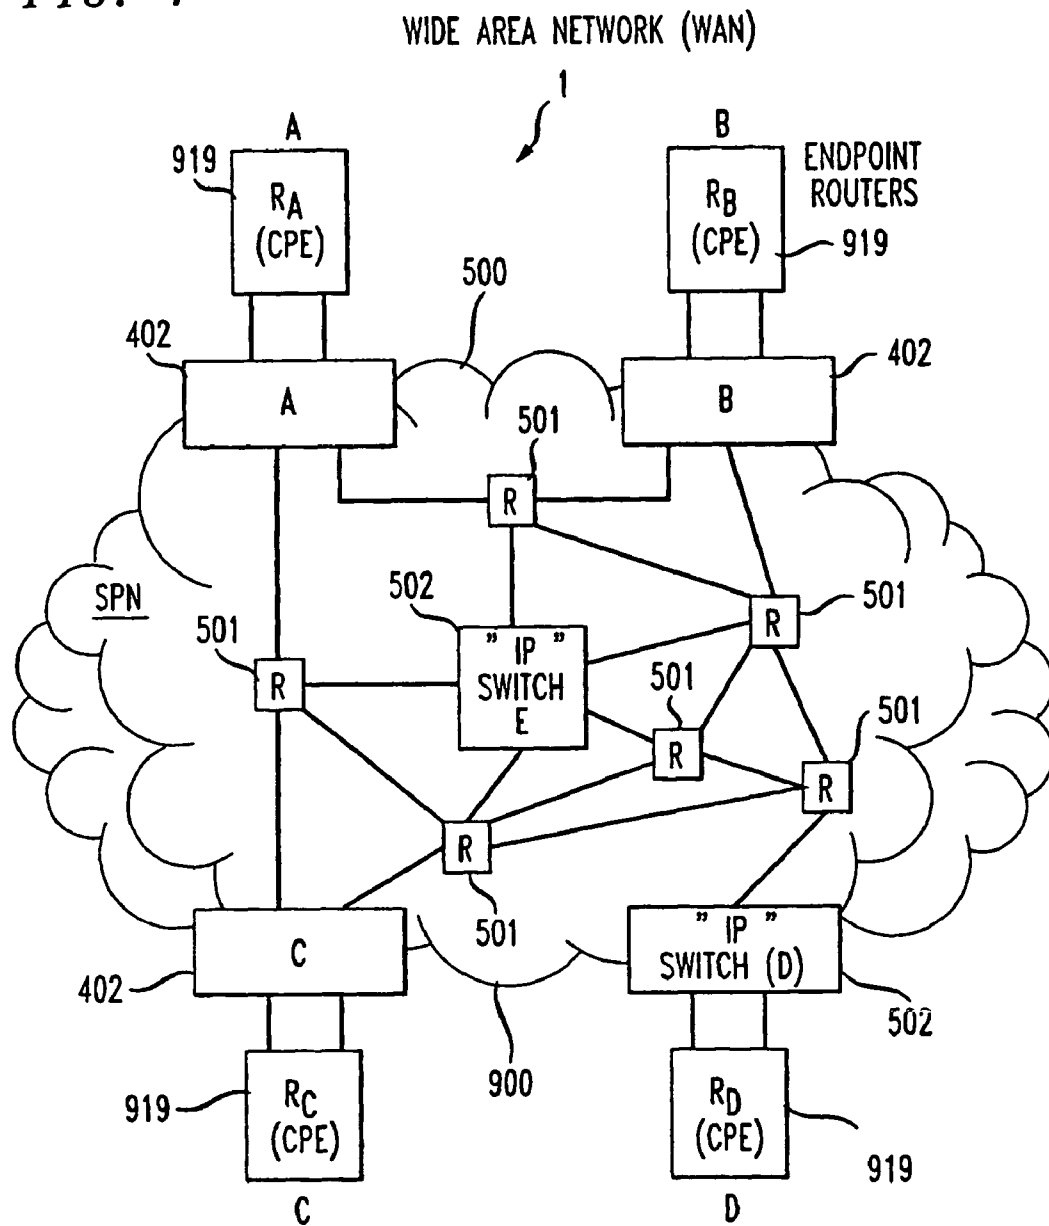
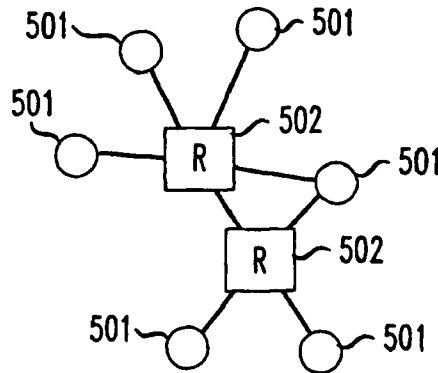
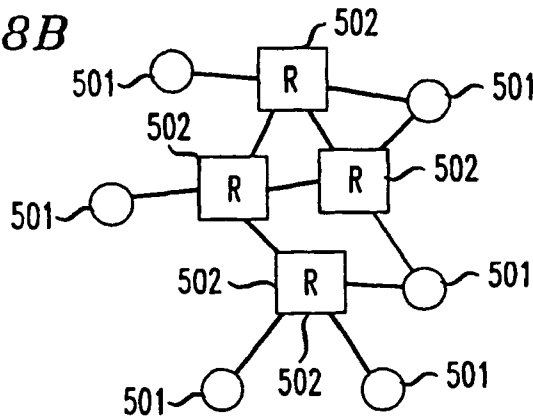


FIG. 8A



INITIALLY, USE THE FR NETWORK TO BACKHAUL TRAFFIC TO THE ROUTER COMPLEXES FOR SWITCHING

FIG. 8B



AS TRAFFIC VOLUME BUILDS, MORE ROUTING COMPLEXES ARE INSTALLED, MEANING LESS BACKHAUL TO THE NEAREST ONE

FIG. 10

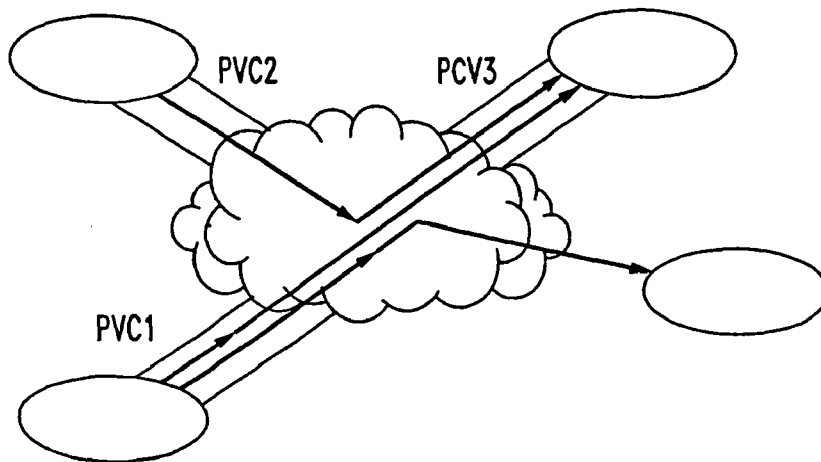


FIG. 9

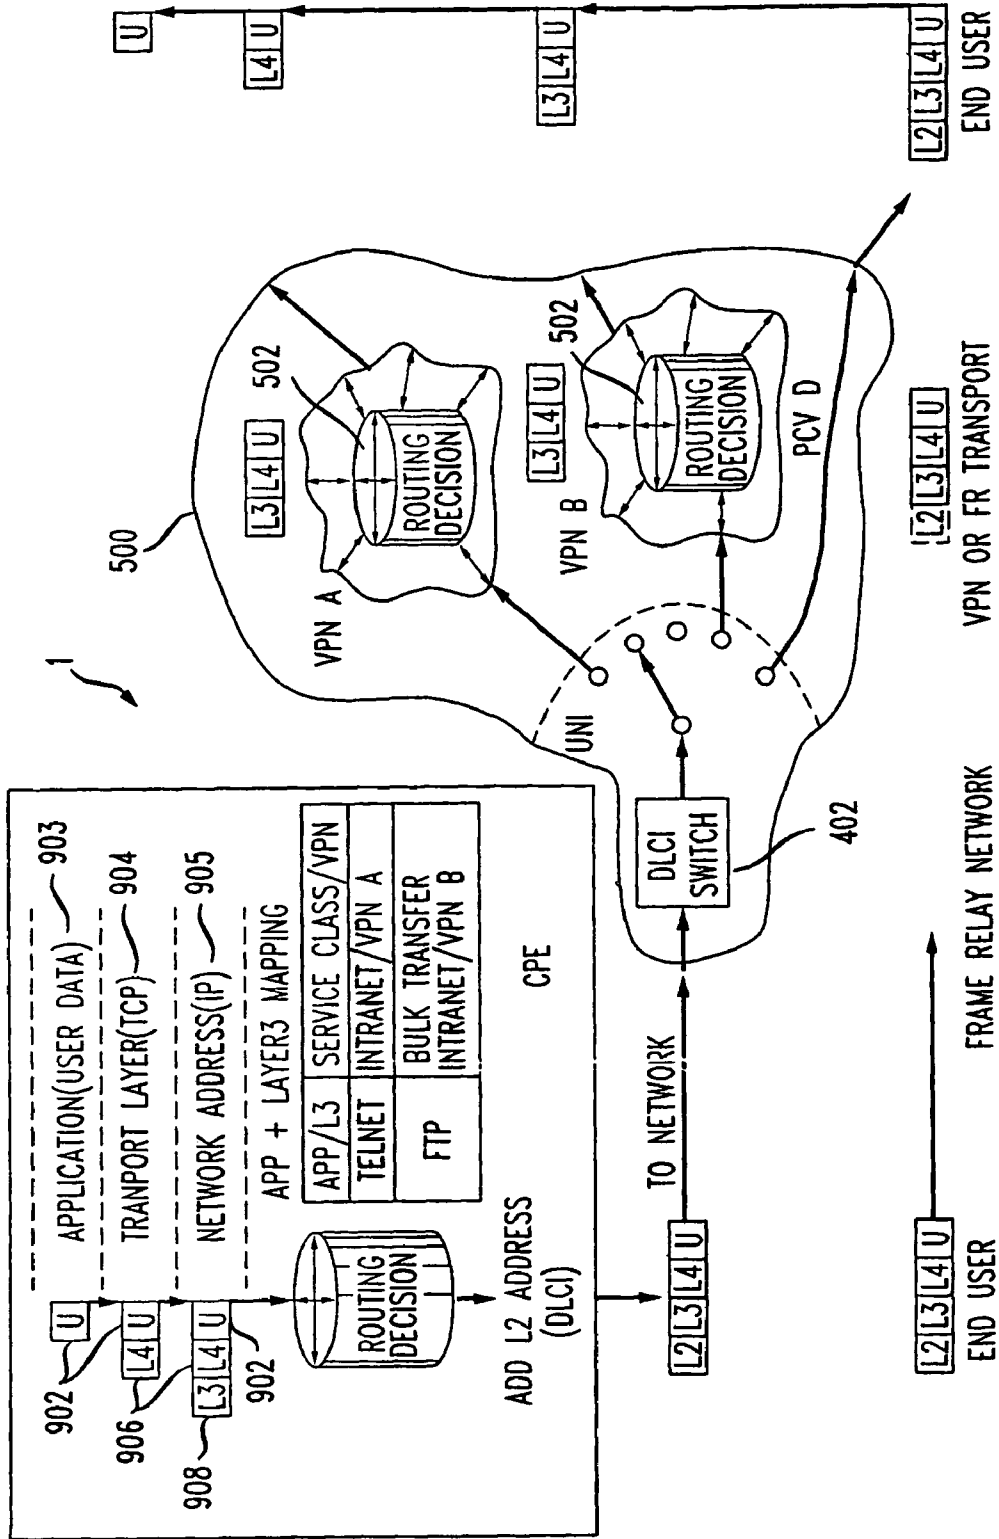
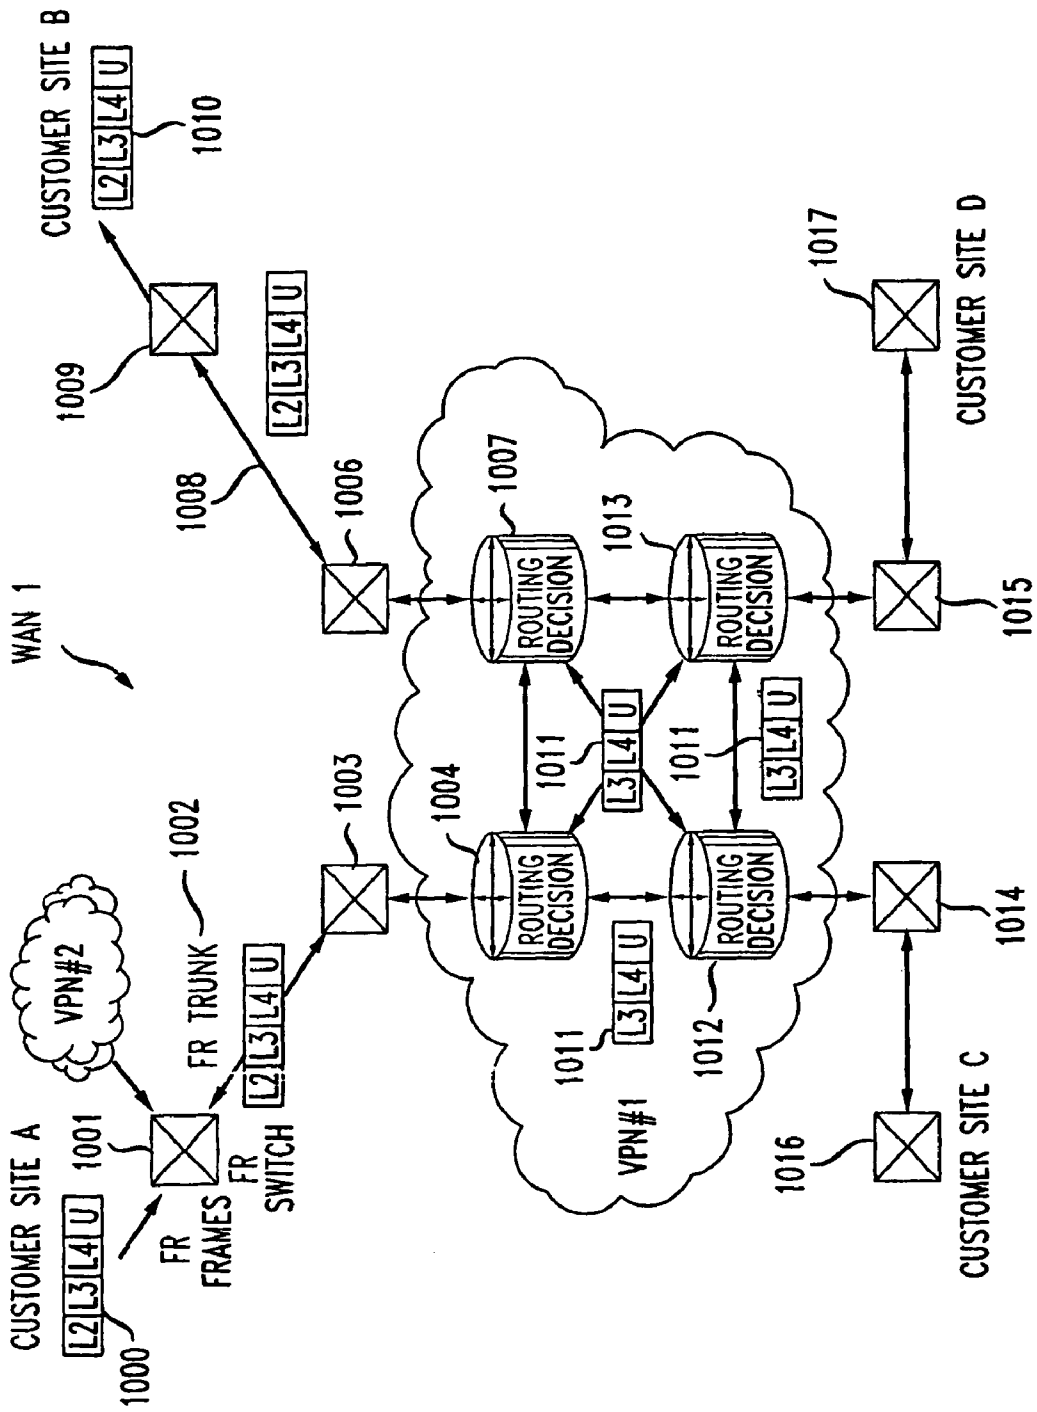


FIG. 12



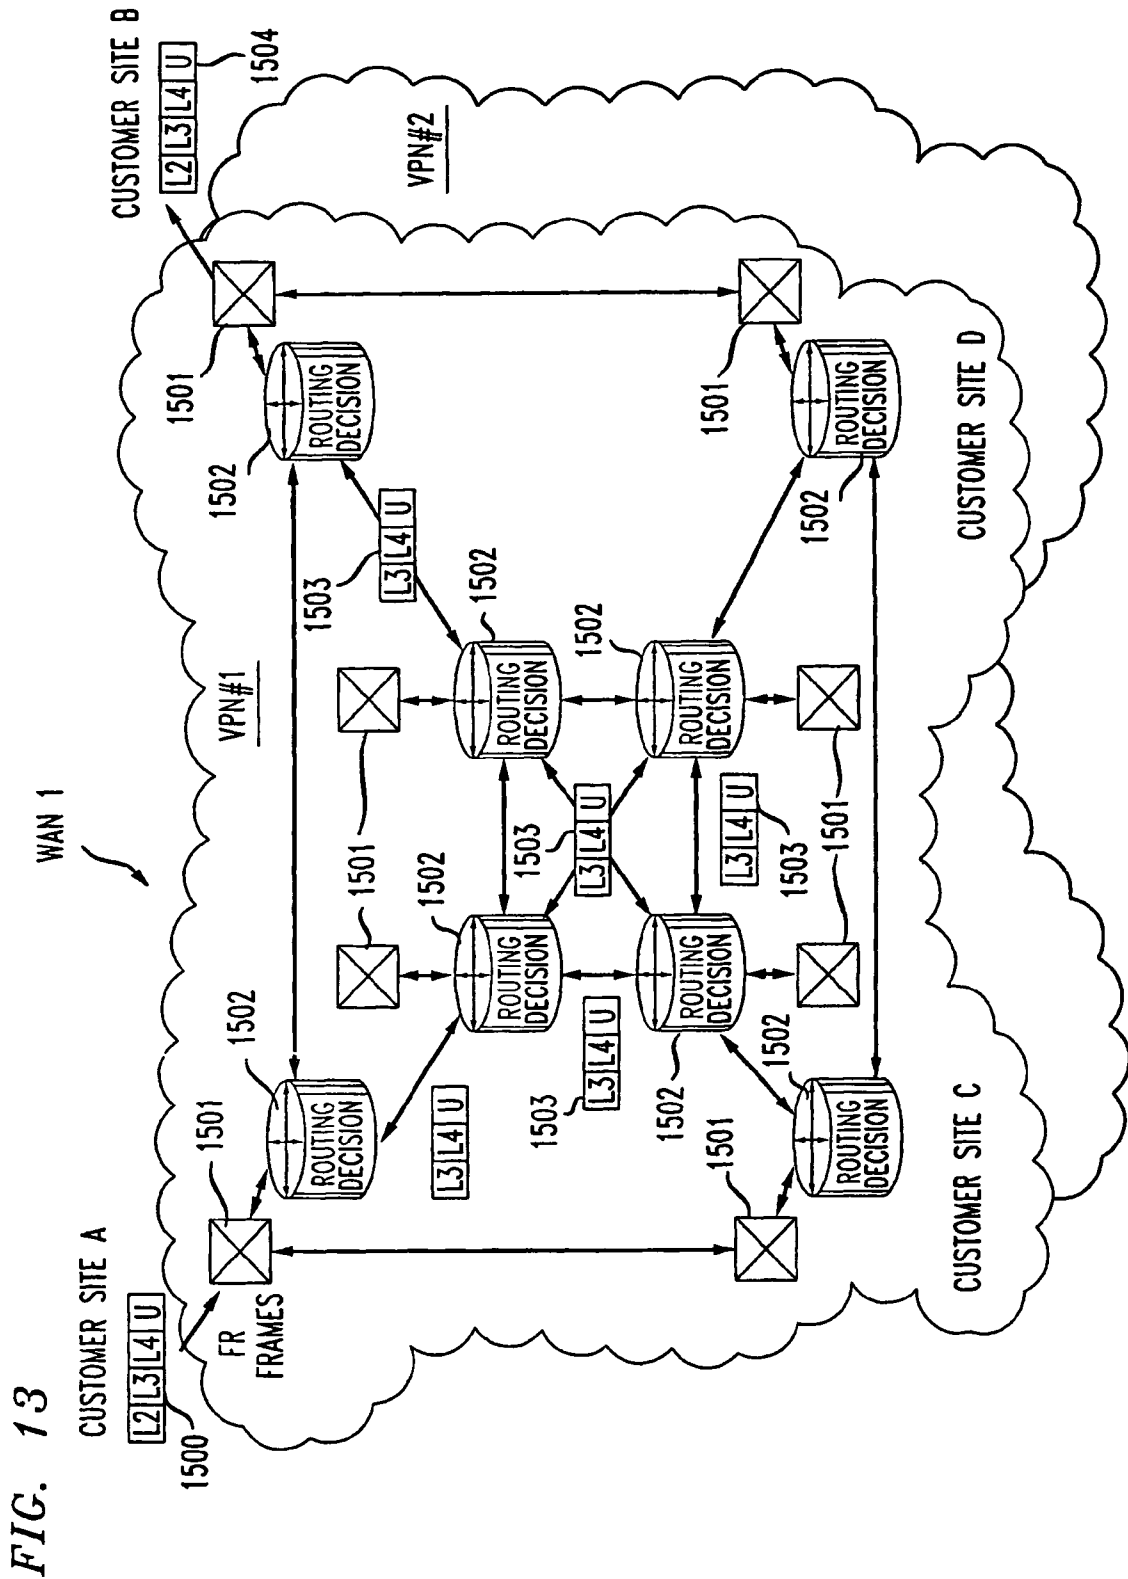
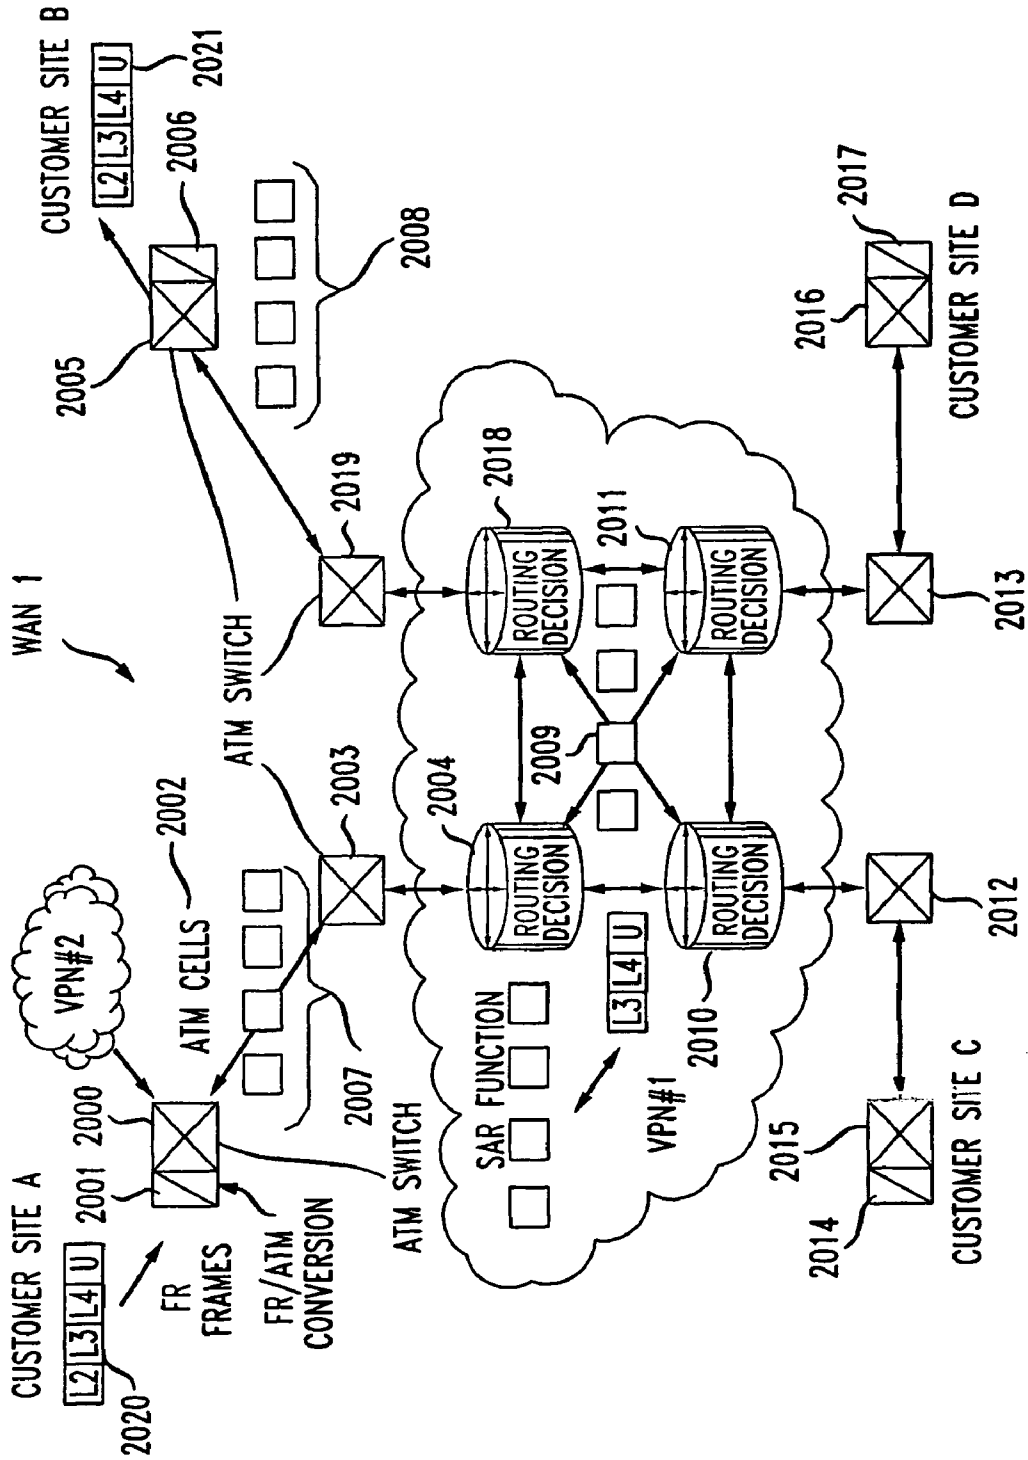


FIG. 14



TRAFFIC MANAGEMENT FOR FRAME RELAY SWITCHED DATA SERVICE

The present application is a continuation of U.S. patent application Ser. No. 09/654,616, filed Sep. 1, 2000 entitled TRAFFIC MANAGEMENT FOR FRAME RELAY SWITCHED DATA SERVICES, now U.S. Pat. No. 6,847,611, issued Jan. 25, 2005, which is a Continuation of U.S. patent application Ser. No. 08/988,424, filed Dec. 10, 1997 now U.S. Pat. No. 6,188,671, issued Feb. 13, 2001 entitled TRAFFIC MANAGEMENT FOR FRAME RELAY SWITCHED DATA SERVICES, which claimed priority from provisional application Ser. No. 60/051,564 entitled FRAME RELAY SWITCHED DATA SERVICE filed on Jul. 3, 1997, and is related by subject matter to U.S. patent application Ser. No. 09/551,399, filed Apr. 17, 2000, entitled FRAME RELAY SWITCHED DATA SERVICE, which is a continuation application of U.S. patent application Ser. No. 08/988,159, filed Dec. 10, 1997, entitled FRAME RELAY SWITCHED DATA SERVICE, now U.S. Pat. No. 6,081,524, by the same inventors, wherein each of the above applications is herein incorporated by reference.

BACKGROUND OF THE INVENTION

1. Technical Field

The present invention is directed to systems and methods for implementing improved network architectures, and more specifically to systems and methods for routing internet protocol (IP) packets using modified frame relay protocols.

2. Description of the Related Arts

Recently, the popularity of large meshed networks has been increasing. However, large-scale highly-meshed networks can be difficult to implement, maintain, and manage using conventional network technologies.

An example of a conventional mesh configuration is shown in FIG. 1. A wide-area network (WAN) **900** includes a plurality of routers R_A, R_B, R_C, R_D , (customer premises equipment (CPE)) respectively disposed at a plurality of end user locations A, B, C, and D and interconnected to a service providers network (SPN) **901** via respective user-network interfaces (UNI) **920-1, -2, . . . , -n**. The user-network interfaces **920** may be variously configured to be, for example, an asynchronous transfer mode (ATM) switch having a frame relay interface to CPE. Connecting the sites together are logical paths called, for example, permanent virtual circuits (PVCs) $P_{A-C}, P_{A-D}, P_{B-D}, P_{A-B}, P_{C-B}$, that are characterized by their endpoints at the UNIs **920-1, 920-2, . . . , 920-n** and a guaranteed bandwidth called the committed information rate (CIR).

FIG. 2 provides a detailed view of the flow of data across the WAN **900**. There exists a plurality of layers of protocol over which communications may occur. For example, the well-known layers of the International Standards Organizations (ISO) Open Systems Interconnect Model having layers from a physical layer (layer 1), a datalink layer (layer 2), a network layer (layer 4), up through and including an application layer (layer 7). Under this model, user data **902** is generated by a user application running at the application layer **903**. At the transport layer (layer 4) **904**, a source and destination port address **906** (as part of the TCP header (layer 4)) may be added to the user data **902**. At the network layer (layer 3) **905**, an additional header (i.e., an IP header (layer 3)) containing source and destination IP addresses **908** may be added. Thus, the layer 3 user data field includes the layer 4 user data **902** plus the layer 4 header **906**. The layer 3 protocol data unit (PDU) **902, 906, 908**, which makes up, for example,

an IP packet **950**, is then passed down to layer 2 **909** in the CPE (routers R_A, R_B, R_C, R_D) that interfaces to the SPN **901**. In the router, a table maps one or more IP addresses (layer 3) **908** to an appropriate PVC or PVCs ($P_{A-C}, P_{A-D}, P_{B-D}, P_{A-B}, P_{C-B}$). The router table is maintained by the customer. Once the correct PVC is located in the routing table, the corresponding data link connection identifier (DLCI) (layer 2) **912** is coded into the header of the frame relay frame **914** (packet). Thereafter, the remainder of the frame relay frame is included and a frame check sum (FCS) is computed. The frame is then passed down to the physical layer and transmitted to the SPN **901**.

At the UNI **920**, the frame is checked for validity to determine if there is a predefined PVC associated with the DLCI **912**. If so, the frame **914** is then forwarded on that PVC through the network along the same path and in the same order as other frames with that DLCI, as depicted in FIG. 2. The layer 2 frame information remains as the packet traverses the frame relay network whether this network is actually implemented as a frame relay network or other network such as an ATM network. The frame is carried to its destination without any further routing decisions being made in the network. The FCS is checked at the egress UNI, and if the frame is not corrupted, it is then output to the UNI associated with the end user.

As is well known in the art, FIGS. 1-3 provide exemplary diagrams of how the frame relay data packets are assembled at the various ISO layers using the example of TCP/IP protocol transport over a frame relay data link layer. The example shows how the user data at the application layer is wrapped in succeeding envelopes, making up the PDUs, as it passes down the protocol stack. Specifically, the composition of the Header field is expanded for detail and is shown in FIG. 5. The data link connection identifier (DLCI) field comprises 10 bits spread over the first and second octet, and allows for 1023 possible addresses, of which some are reserved for specific uses by the standards. As shown in FIG. 3, the DLCI is added to the frame relay header according to what destination IP address is specified in the IP packet. This decision about what DLCI is chosen is made by the CPE, usually a router, based on configuration information provided by the customer that provides a mapping of IP addresses into the PVCs that connect the current location with others across the WAN **900**.

In conventional frame relay, a layer 2 Q.922 frame carries the layer 3 customer data packet across the network in a permanent virtual circuit (PVC) which is identified by a data link connection identifier (DLCI). Thus, the DLCIs are used by the customer as addresses that select the proper PVC to carry the data to the desired destination. The customer data packet is carried across the network transparently and its contents is never examined by the network.

The conventional meshed frame relay network discussed above has a number of limitations. For example, every time a new end user location is added to the meshed network, a new connection is required to be added to every other end user location. Consequently, all of the routing tables must be updated at every end user location. Thus, a ripple effect propagates across the entire network whenever there is a change in the network topology. For large networks with thousands of end user locations, this ripple effect creates a large burden on both the network provider to supply enough permanent virtual circuits (PVCs) and on the network customers in updating all of their routing tables. Further, most routers are limited to peering with a maximum of 10 other routers which makes this network topology difficult to implement. As networks grow in size, the number of PVCs customers need to manage and map to DLCIs increases. Further

complicating the problem is a trend toward increasing meshedness of networks, meaning more sites are directly connected to each other. The result is a growth in the number and mesh of PVCs in networks that does not scale well with current network technologies.

A possible solution for handling large meshed networks is to use a virtual private network (VPN) which interconnects end user locations using encrypted traffic sent via tunneling over the internet. However, VPNs are not widely supported by internet service providers (ISPs), have erratic information rates, and present a number of security concerns.

Another possible solution is the use of frame relay based switched virtual circuits (SVCs). While PVCs (discussed above) are usually defined on a subscription basis and are analogous to leased lines, SVCs are temporary, defined on an as-needed basis, and are analogous to telephone calls. However, SVCs require continuous communications between all routers in the system to coordinate the SVCs. Further, because the tables mapping IP addresses to SVC addresses are typically manually maintained, SVCs are often impractical for large highly-meshed networks. Security is a major concern for SVC networks where tables are mismanaged or the network is spoofed. Further, frame SVCs are difficult to interwork with asynchronous transfer mode (ATM) SVCs.

None of the above solutions adequately address the growing demand for large mesh networks. Accordingly, there is a need for network architectures which enable implementation of large mesh networks having security, low maintenance costs, efficient operations, and scalability.

SUMMARY OF THE INVENTION

Aspects of the present invention solve one or more of the above-stated problems and/or provide improved systems and methods for implementing a network architecture.

A new type of data transport service takes advantage of the existing base of frame relay customer premises equipment (CPE) and customers while offering a new mechanism for providing extensible service features to those customers. In the new service, data link connection identifiers (DLCIs) may be used by the CPE to select among service types, feature sets, and closed user groups (CUGs). The DLCI is used in the layer 2 frame that conveys the user data to the network. The layer 3 user data packet is extracted from the layer 2 frame and the layer 3 address information for the (routable) protocol is used to route the user data packet over a high-performance packet switched network, according to the service class/feature set selected by the DLCI. At the destination, the layer 3 data packet is again enclosed in a layer 2 frame with a DLCI that indicates to which service group it belongs. The frame is then forwarded to the CPE. Use of this technique will allow the existing frame relay CPE to support, over the same physical interface, conventional frame relay service with a range of DLCIs that are linked to logical paths such as permanent virtual circuit (PVCs), as well as a range of DLCIs that are linked to service and/or feature sets. This will allow a robust method for extension of new services to the frame relay installed base, with minimal impact to existing customer equipment.

In some aspects of the invention, frame relay DLCIs are used for selecting among various service categories. This differs significantly from conventional frame relay, which uses DLCIs only to select PVCs and/or switched virtual circuits (SVCs). Service categories may include, but are not limited to, communication via the public internet, communication via a local intranet, communication within a closed user group (CUG), communication with an extranet (e.g., a

network of trusted suppliers or corporate trading partners), live audio/video transmission, multicasting, telephony over internet protocol (IP), or any combination thereof. Thus, the concept of a frame relay PVC is significantly expanded by aspects of the present invention. For example, the location of an intended network endpoint recipient is not necessarily determined by a DLCI at a sending network endpoint. The DLCI may represent a service category with the intended recipient indicated by an IP address within the frame relay packet. This results in a significant benefit to network customers because, unlike that of conventional frame relay, customers no longer need to update their local DLCI tables each time a network customer with whom they wish to communicate is added or removed from the network. Thus, the customer's burden of network administration is substantially reduced.

In sub-aspects of the invention, some DLCIs may be used to select among service categories (service category DLCIs) while in the same network other DLCIs may be used to select conventional PVCs and/or SVCs (conventional DLCIs). In other words, conventional frame relay may be mixed with aspects of the present invention within the same network, allowing aspects of the present invention to be incrementally implemented in existing conventional frame relay networks.

In further aspects of the invention, addressing contained in multiple layers (e.g., as defined by the Open System Interconnection model) are compared with each other in a network to determine routing errors. If the addressing in the layers are consistent with each other, then the associated data is routed without interruption. On the other hand, if the addressing in the layers is inconsistent with each other, the associated data may be specially handled. For example, the data may be discarded, sent to a pre-determined address, and/or returned to the sender. This address comparison may be applied to the sending address and/or the destination address. An advantage of this multiple layer address comparison is that network security is increased. For instance, problems such as spoofing, which is the practice of purposely providing an incorrect sending internet protocol (IP) address, are better controlled by such a method.

In still further aspects of the invention, routing look-up tables within the network are separated such that, for example, each customer, closed user group (CUG), extranet, and/or intranet may have its own private partition and/or separate table. This can provide greater network speed because a router need not scan the entire available address space for all network customers at once. Furthermore, data security is improved because the risk of sending data to a wrong recipient is reduced.

In yet further aspects of the invention, layer 3 and/or layer 4 IP address information is utilized to route the fast packets through the network.

In even further aspects of the invention, new network traffic management techniques and measurements are defined. For example, in some traffic-management aspects of the invention, committed delivery rates (CDRs) may be assigned to one or more UNIs. A CDR is the average minimum data rate that is guaranteed to be delivered to a given UNI when sufficient traffic is being sent to the UNI. In further traffic-management aspects of the invention, a destination rate share (DRS) is assigned to one or more UNIs. The DRS may be used to determine the share of traffic that a given UNI may send through the network. If several UNIs are simultaneously offering to send traffic to the same destination UNI, then each sending UNIs share of the network may be determined by its own DRS and the DRSs of the other sending UNIs.

These and other features of the invention will be apparent upon consideration of the following detailed description of preferred embodiments. Although the invention has been defined using the appended claims, these claims are exemplary in that the invention is intended to include the elements and steps described herein in any combination or subcombination. Accordingly, there are any number of alternative combinations for defining the invention, which incorporate one or more elements from the specification, including the description, claims, and drawings, in various combinations or sub-combinations. It will be apparent to those skilled in network theory and design, in light of the present specification, that alternate combinations of aspects of the invention, either alone or in combination with one or more elements or steps defined herein, may be utilized as modifications or alterations of the invention or as part of the invention. It is intended that the written description of the invention contained herein covers all such modifications and alterations.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing summary of the invention, as well as the following detailed description of preferred embodiments, is better understood when read in conjunction with the accompanying drawings. For the purpose of illustration, embodiments showing one or more aspects of the invention are shown in the drawings. These exemplary embodiments, however, are not intended to limit the invention solely thereto.

FIG. 1 illustrates a wide area network (WAN) having routers as CPEs and PVCs between customer locations.

FIG. 2 shows data flow through the WAN shown in FIG. 1.

FIGS. 3-5 show the construction and flow of data packets through the network.

FIG. 6 shows a block diagram of a network architecture in accordance with aspects of the present invention.

FIG. 7 shows a detailed block diagram of the network illustrated in FIG. 6.

FIG. 8A-8B shows a migration path for incorporating aspects of the invention into conventional network architectures.

FIG. 9 shows data flow through the network architecture of FIG. 6.

FIG. 10 shows application based prioritization through the network architecture of FIG. 6.

FIG. 11 illustrates an exemplary embodiment of a means to apportion services through the network of FIG. 6.

FIGS. 12-14 illustrate data flow through exemplary WANs 1.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Exemplary embodiments of the present invention allow the large installed base of frame relay customer premises equipment (CPE) to be maintained by using the same interface in a different way to deliver new sets of services and features to the customer. For example, the data link connection identifier (DLCI) known from the frame relay protocol may be used to select among several virtual private networks with differing address spaces, feature sets, and/or conventional permanent virtual circuits (PVCs).

Referring to FIG. 7, a block diagram of a wide area network (WAN) 1 incorporating aspects of the present invention is shown. The WAN 1 includes a plurality of customer premise equipment (CPE) system, for example routers located at each of the end user locations and interconnected via one or more service providers networks (SPNs) 500. The SPN 500 is

typically connected to a plurality of endpoint routers 919 via a plurality of corresponding user network interfaces (UNIs) 402 and/or one or more internet protocol (IP) switches 502. The IP switches 502, UNIs 402, and/or routers/switches 501 may be interconnected so as to form a meshed network (e.g., a partial or fully meshed network). Additionally, the wide area network (WAN) 1 may contain any number of IP switches 502 located within the WAN 1 such that it is not connected directly to any endpoint routers 919, and/or one or more IP switches 502 may be located at an interface between the SPN 500 and an endpoint router 919. In further embodiments of the invention, there may be multiple endpoint routers 919 associated with a UNI 402/IP switch 502 and/or multiple UNIs 402/IP switches 502 associated with an endpoint router 919.

The network architecture of the WAN 1 allows the number of IP switches to increase as customers are transitioned to the new service. For example, as shown in FIG. 8A, initially there may be only a small number (e.g., one, two, three, etc.) of IP switches installed in the system. Where only a small number of IP switches are included in the network, traffic originating from non-IP enabled UNIs 402 (e.g., UNI A) may be routed to an IP switch 502 elsewhere in the network. Although this creates some negligible inefficiencies in backtracking it nonetheless allows a migration path to the new network architecture without simultaneously replacing all routers 501. However, as more and more users are transitioned to the new network architecture of WAN 1, more and more IP switches can be added (FIG. 8B) to accommodate the increased load. In many embodiments, it may be desirable to eventually convert each UNI 402 to an IP switch 502 such that IP routing may be accomplished at the edge of the network.

In some embodiments, the WAN 1 may include a combination of conventional network switches and/or routers 501 in addition to IP switches 502. On the other hand, every switch in the SPN 500 may be an IP switch 502. Alternatively, the WAN 1 may contain only a single IP switch 502. The IP switches 502 may be variously configured to include a suitable multi-layer routing switch such as a Tag Switch from Cisco. Multi layer routing switches may also be utilized from vendors such as Ipsilon, Toshiba, IBM, and/or Telecom. IP switches are currently being developed to replace endpoint routers so that customer premise equipment (e.g., Ethernet local area network (LAN) equipment) can connect directly to an asynchronous transfer mode (ATM) network. Aspects of the present invention propose using IP switches in a different manner to maintain the huge installed base of customer premise equipment while avoiding the limitations of previous systems. Accordingly, the IP switches in accordance with embodiments of the invention are disposed within the SPN 500 and modified to provide suitable routing and interface functions.

In some embodiments of the invention, an IP switch 502 acts as a multi-layer switch. For example, an IP switch 502 may receive ATM cells, switching some or all of the ATM cells based upon the content of IP packets encapsulated within the ATM cells. Thus, IP addressing may be used by an IP switch 502 to determine an ATM virtual path for sending ATM cells to a destination UNI 402. In further embodiments of the invention, higher layer addressing (e.g., transmission control program (TCP) logical ports at layer 4) may also be used by an IP switch 502 as a basis for switching ATM cells to provide a path through the SPN 500. In still further embodiments of the invention, an IP switch 502 uses IP addresses and/or TCP logical ports to make quality of service (QOS) decisions.

In further embodiments of the invention, an endpoint router **919** may encapsulate one or more IP packets in frame relay frame **914**. In this event, the frame relay frames may be transmitted between an endpoint router **919** and a corresponding UNI **402** and/or IP switch **502**. The endpoint router **919** encapsulates IP packets **950** with frame relay frames **914**. Further, the endpoint router **919** may set the DLCI of each frame relay frame **914** according to a particular service category (if a service category DLCI is used) that the user has selected. For example, the various service categories may include the public internet, communication via a local intranet, communication within a closed user group (CUG), communication with an extranet (e.g., a network of trusted suppliers or corporate trading partners), live audio/video transmission, multicasting, telephony over internet protocol (IP), or any combination thereof. Thus, the concept of a frame relay PVC is significantly expanded by aspects of the present invention. For example, the location of an intended network endpoint recipient is not necessarily determined by a DLCI at the endpoint routers **919**.

In further embodiments of the invention, a UNI **402** may receive frame relay frames **914** from an endpoint router **919** and divides and encapsulates frame relay frames into, for example, smaller fixed-length ATM cells. The UNI **402** may further translate the frame relay DLCI into an ATM address (e.g., a virtual path identifier/virtual channel identifier (VPI/VCI)). There are various methods which may be used to translate DLCIs to VPI/VCIs. For example, the Network Interworking Standard as defined in Implementation Agreement #5 of the Frame Relay Forum, and/or the Service Interworking Standard as defined in Implementation Agreement #8 of the Frame Relay Forum may be utilized. An ATM address associated with a service category DLCI defines an ATM virtual path via network routers to an IP switch **502**. Thus, ATM data associated with a service category DLCI is ultimately sent to an IP switch **502**. However, ATM data associated with a conventional DLCI may or may not be sent to an IP switch **502** and may be routed through the network without passing through an IP switch **502**. Thus, both translated IP data and conventional PVC data may be present in the SPN **500** and/or WAN **1**.

In further embodiments of the invention, a UNI **402** and/or a network router **501** may send data to a predetermined IP switch **502**. In even further embodiments of the invention, a UNI **402** and/or a network router **501** selects which IP switch **502** to send data to based upon an algorithm (e.g., based on network traffic flows, the relative distance/location of an IP switch **502**, the type of data being sent, and/or the service category selected). In still further embodiments of the invention, a UNI **402**, network router **501**, and/or IP switch **502** may send the same data to more than one UNI **402**, network router **501**, and/or IP switch **502**, depending upon, for example, a service category or categories.

In further embodiments of the invention, a UNI **402**, an IP switch **502**, and/or a network router **501** compares an ATM VPI/VCI **303-305** address with an IP address for the same data. If the two addresses are inconsistent, then the ATM cell may be discarded, sent to a pre-determined address, and/or returned to the sending location. In even further embodiments of the invention, layers above the layer 3 IP layer may be used for address and/or service class generation/discrimination. For example layer 4 of the ISO addressing scheme and/or other application level data may be utilized to determine particular service classes.

Referring specifically to FIG. 9, the path of user data flowing through an exemplary WAN **1** is shown. As in the frame relay case, user data at the application layer and layer 4

requires the addition of a layer 3 network address header. In the CPE a decision is made based on information in layers 3 and 4 about which virtual private network (VPN), service class, or conventional PVC the packet should be routed to. Thus, a packet with layer 4 information indicating it is a telnet (interactive) application and layer 3 information that it is an internal company address might go to VPN A for a low-delay intranet class of service. Another packet that is part of a file transfer protocol (FTP) file transfer might go to VPN B with a lower service class, and a third packet going between two heavily utilized applications might go on a dedicated PVC D. These decisions are coded as different DLCI values, inserted in the layer 2 frame, and sent into the UNI.

At the UNI A **402**, the switching based on the DLCI takes place. The packet may be routed to IP switch **502** in the center of the SPN **500**. The first packet has its layer 2 frame stripped off as it is forwarded to VPN A. Within VPN A, the layer 3 address is now used to make routing decisions that send the packet to its destination UNI. Thus, no PVC need be established ahead of time for that path, and conventional routing methods and protocols can be used, as well as newer short-cut routing techniques. This permits VPN A to provide a high mesh of connectivity between sites without requiring the customer to configure and maintain the mesh as a large number of PVCs. The packet forwarded to VPN B is treated similarly except that VPN B is implemented with a lower service class (e.g. higher delay). Finally, the packet forwarded to PVC D has its layer 2 frame intact and passes through the network as a conventional frame relay frame. This allows customers to maintain their current connectivity of PVCs for their high utilization traffic paths, but still have a high mesh of connectivity through various VPNs.

Thus, in various aspects of the invention, the WAN **1** and/or SPN **500** may be any suitable fast packet network receiving frame relay data packets having user data in a user data field. The WAN **1** and/or SPN **500** then switches packets using one or more IP switches **502** responsive to the user data. The user data may be used to discriminate between a plurality of different service categories based on the user data. Routing over the WAN **1** and/or SPN **500** may be responsive to at least one of the different service categories including discriminating based on multicast data. Additionally, the WAN may generate a fast packet address field responsive to the IP packet data and route the IP packet through the fast packet network responsive to the fast packet address field. Further, layer 4 information may be utilized to determine the quality of service. The quality of service may include, for example, one or more of the following: an information rate, priority information, delay, loss, availability, etc. Security features may be implemented in the IP switch such that routing tables for each of the users are separated based on one or more service categories and/or users. In this manner the system is made more secure. Still further, the system may receive a plurality of frame relay packets over a permanent virtual circuit (PVC) at a first node in an asynchronous transfer mode (ATM) network, generate an ATM address based on a data field other than a data link connection identifier (DLCI) within the frame relay packets, and then route the packets through the ATM network based on the ATM address. The routing of packets may be responsive to one of a plurality of service categories. The system may provide separate routing tables within an ATM switch for each of a plurality of different service categories. The different service categories may be determined using internet protocol (IP) data within a data field of a packet passed by the ATM switch. In a fast packet network, a fast packet switch may compare an address of a fast packet with a layer 3 internet protocol (IP) address contained within the fast packet and

determining whether the fast packet address is consistent with the layer 3 IP address. Further, for security, hardware circuits and/or software may be provided for examination of a sending address or a destination address. Further, packets may be discarded responsive to an inconsistency being detected. The WAN 1 may include customer premises equipment (CPE) and an asynchronous transfer mode (ATM) switch coupled to and receiving from the CPE frame relay data packets, and including address translation circuitry for translating data link connection identifiers from the frame relay data packets into ATM addresses representing a plurality of virtual private networks based on a predetermined service category associated with a particular DLCI; or the WAN 1 may include customer premises equipment (CPE) and a fast packet switch coupled to the CPE via one or more permanent virtual circuits and receiving frame relay data packets, the fast packet switch including address translation circuitry for translating user data within the frame relay data packets into fast packet addresses.

In embodiments of the present invention, data security is enhanced in that data may be easily and accurately checked for inconsistencies at the destination. This is because these embodiments operate using both layer 2 and layer 3 addressing information. As an illustration, assume that a frame relay frame having a DLCI indicating VPN 1 (e.g., the corporate intranet) arrives in a network switch/router with an IP address of a particular corporate accounting system. However, since the VPN processor has available to it the DLCI of the packet (and thus information about the source of the packet), the VPN processor may cross-check the DLCI with the source IP address in the packet to see if the source IP address is in the range known from the originating site. Thus, the problem associated with the spoofing of IP source addresses may be significantly reduced.

In still further embodiments of the invention, a UNI 402, an IP switch 502, and/or a network router 501 has separate and/or partitioned routing look-up tables. Routing tables may be separated based upon service category, customer or user, and/or UNI 402. Thus, in some embodiments, within a VPN, a customer or user may have an individual routing table containing the customers IP network address information. In some embodiments, since the DLCI identifies the source of a frame, the DLCI may be used as an index by an IP switch, network router, and/or UNI for determining which routing table to use. This allows customers to have their routing table size and speed governed by their individual address space, thus speeding the routing process considerably. The use of separate routing tables also provides an added measure of security, as packets cannot be mis-routed due to errors or updates in routing information related to other customers.

In some embodiments, a router has multiple data space images paired with a single instruction space image of the routing software. Thus, for example, as packets arrive from Customer A, the routing software uses the data image for a routing table associated with Customer A to make a routing decision. In further embodiments, a single software image is used, but additional indices corresponding to customers are added to the routing tables. In still further embodiments, instruction execution and data handling are processed separately. This may be accomplished by the use of separate processors, one for instruction execution and one for data handling.

FIG. 12 illustrates an exemplary WAN 1 having both conventional routers and IP switches incorporating aspects of the invention. In this exemplary WAN 1, a routing element 1004 and switch 1003 are connected to Customer Site A via frame relay switch 1001. Routing element 1007 and switch 1006 are

connected to Customer Site B via frame relay switch 1009. Routing element 1012 and switch 1014 are connected to Customer Site C via frame relay switch 1016. Routing element 1013 and switch 1015 are connected to Customer Site D via frame relay switch 1017. In this exemplary WAN 1, incoming frames 1000 from Customer Site A may be encoded with a layer 2 DLCI specifying VPN #1 as the layer 2 destination and a layer 3 address pointing to Customer Site B. In such a case, frame relay switch 1001 switches the frames over a frame relay trunk 1002 to switch 1003 which has layer 3 routing element 1004 associated with it. After the frame is received by switch 1003, the frame is forwarded to router 1004 which implements short-cut routing as described above. The router/switch 1003, 1004 uses the layer 2 information to discriminate between different source customers. The layer 2 information may then be discarded. Next, the layer 3 information in combination with a routing table is used to make a routing decision. In this case, the routing decision would result in a layer 3 PDU 1011 being forwarded to router/switch 1006, 1007. The layer 3 PDU 1011 is then encapsulated with a layer 2 frame, the frame in this case being addressed to Customer Site B. Switch 1006 then forwards the frame via a trunk 1008 to frame relay switch 1009. At the egress port of frame relay switch 1009, the DLCI of frame relay frame 1010 is replaced with a value indicating that the frame originated from, in this case, VPN #1. The frame relay frame 1010 is then delivered to the Customer B router.

As the service grows, the functionality for making the VPN routing decisions may be migrated closer to the customer and may eventually be present in every switching node, as shown in FIG. 13. This can reduce the backhaul previously needed to get to the router/switch processing nodes and allow for optimal routing using all the nodes in the WAN 1 and/or SPN 500. In the exemplary embodiment of FIG. 13, VPN #1 is connected to Customer Sites A, B, C, and D. Here, every switching node includes a switch 1501 and a routing element 1502. frame relay frames 1500 having a DLCI directed to Customer Site B may be sent from Customer Site A. In such a case, frames 1503 would be sent through VPN #1 via switching nodes 1501, 1502, and frames 1504 would be received at Customer Site B.

In some embodiments, an ATM core network may be used for data transport, and frame relay interfaces may be used to interface with the customer. An exemplary embodiment using an ATM core network is shown in FIG. 14. In this embodiment, switch 2003 and router 2004 are connected to Customer Site A via switch 2000 and a frame relay/ATM conversion unit 2001. Switch 2019 and router 2018 are connected to Customer Site B via switch 2005 and frame relay/ATM conversion unit 2006. Switch 2012 and router 2010 are connected to Customer Site C via switch 2015 and frame relay/ATM conversion unit 2014. Switch 2013 and router 2011 are connected to Customer Site D via switch 2016 and frame relay/ATM conversion unit 2017. Assuming that Customer Site A is sending frames 2020 destined for Customer Site B, incoming layer 2 frames may be encapsulated for transport into ATM cells at switch 2000 according to, for example, the Network Interworking Standard. Such encapsulation may, for example, occur in conversion unit 2001, external to ATM switch 2000. ATM cells 2002 may be sent down an ATM PVC designated for VPN #1 processing. ATM cells 2002 may then be forwarded to switch 2003 and router/switch 2004 (which may be attached to switch 2003), where the ATM cells may be reassembled to obtain the layer 3 packet information for routing within VPN #1. Once the address information has been extracted from the layer 3 packet, the packet may be segmented again into ATM cells 2009 that can be transferred

through the network. After being sent through router/switch **2018**, **2019**, ATM cells **2008** may be converted from cells to frames at the external conversion unit **2006** and switch **2005**. Customer Site B would then receive frame relay frames **2021**. Thus, an extra segmentation and reassembly (SAR) cycle may be required when using an ATM backbone with a core of router/switches. However, if the VPN processing is pushed outward to edge switches, the extra SAR cycle may be eliminated. The extra SAR cycle may be eliminated because conversion from frame relay frames to ATM cells may take place in the same unit where VPN routing decisions are made.

Traffic management may be variously configured in the WAN **1** and/or the SPN **500**. For example, from a customer's viewpoint, the WAN **1** and/or SPN **500** may ensure certain traffic rates for the customer.

In a network, data traffic may be sent from multiple sources to a single destination (multi-point to point). A source is defined as the user transmitting side of, for example, a UNI (i.e., the customer side of a UNI, which may be external to a WAN and/or to a VPN), a switch, an IP switch, and/or a router at or near the edge of a network. A destination is defined as the user receiving side of, for example, a UNI (i.e., the network side of a UNI), a switch, an IP switch, and/or router at or near the edge of a network. Traffic that is offered for transmission by a source to the WAN **1** and/or SPN **500** is defined as the offered traffic. Further, a VPN source and a VPN destination are a source and destination, respectively, which belong to a given VPN. A given UNI, if simultaneously sending and receiving, may simultaneously be a source and a destination. Furthermore, a given source may offer data traffic to multiple destinations, and a given destination may receive traffic from multiple sources.

In some embodiments of the invention, a committed delivery rate (CDR) may be assigned to each destination. The CDR is defined as the average number of bits per second that the WAN **1** and/or SPN **500** is committed to deliver to a given destination, wherein the average may be calculated over a fixed or variable time window. Although the word average will be used throughout, any other similar algorithm may be used, such as the mean, the sum, or any other useful measurement and/or statistical calculation. If the average rate of aggregate offered traffic (i.e. the total offered traffic) from one or more sources to a given destination is greater than or equal to a given destination's assigned CDR, then the WAN **1** and/or SPN **500** may guarantee to deliver traffic addressed to the destination at an average rate equal to or greater than the CDR. If the average rate of aggregate offered traffic is less than the CDR, then the WAN **1** and/or SPN **500** may deliver the offered traffic to the destination at the aggregate offered traffic rate (100% of the offered traffic). To clarify, let the number of active sources sending traffic to a particular destination be N. As will be described in more detail below, a source may be considered active during a given time window if the source offers at least a threshold amount of traffic to the WAN **1** and/or SPN **500** within the given time window. Let S_i be the average offered traffic rate, or offering rate, from each source i toward a single given destination, wherein $i=[1, \dots, N]$. Further, let R be the total rate at which the WAN **1** and/or SPN **500** actually delivers traffic to the destination. Then, the WAN **1** and/or SPN **500** will provide that:

$$R \geq CDR \quad \text{if} \quad \sum_i S_i \geq CDR$$

$$R = \sum_i S_i \quad \text{otherwi.}$$

If the aggregate offered traffic rate S_i does not exceed the CDR, then 100% of the offered traffic from each source i may

be delivered through the WAN **1** and/or SPN **500** to the destination. However, when the aggregate offered traffic rate S_i exceeds the CDR, the WAN **1** and/or SPN **500** may have the discretion to throttle back or reduce the delivery rate of offered traffic from some or all of the active sources. Delivery may be reduced by an amount such that the total rate of traffic delivery R to a destination is at least equal to the destination's assigned CDR. In the situation where R is reduced by the network, it may be desirable to enforce fairness for each source. In other words, it may be desirable to ensure that no single source may be allowed to be greedy by obtaining a disproportionate amount of network bandwidth at the expense of other sources.

To provide for fair access to the WAN **1** and/or SPN **500**, in some embodiments each source is assigned at least one destination rate share (DRS). A DRS is a rate, measured in data units per unit of time (e.g., bits per second). A separate DRS and/or set of DRSs may be assigned to each source and/or group of sources. Further, the DRS or DRSs for a given source may depend upon the destination or set of destinations that the source may send traffic to. In other words, each source i may be assigned at least one DRS, corresponding to the DRS assigned between a source i and a given destination (or set of destinations). Thus, in some embodiments, the DRS may be different for a given source depending upon which destination it is sending traffic to. In further embodiments, the DRS for a given source may be constant, independent of the destination.

When a source i offers traffic at an average rate S_i exceeding the CDR of a particular destination, fairness may be achieved by ensuring that each source is allowed to transmit at least its fair share of the CDR. A source's fair share of the destination's CDR is defined as the source's DRS divided by the aggregate DRS of active sources transmitting to a given destination. Thus, each source's fair share, r_i , of the CDR may be defined as the following:

The actual network

$$r_i = \frac{DRS_i}{\sum_i DRS_i} CDR$$

transmission rate, T_i , that the WAN **1** and/or SPN **500** chooses as conforming traffic guaranteed to be delivered from each source to a given destination may satisfy the following:

$$\text{when} \quad \sum_i S_i \geq CDR$$

$$T_i \geq \min(r_i, S_i)$$

Thus, in these embodiments the WAN **1** and/or SPN **500** may enforce fairness by reducing one or more sources' actual network transmission rate T_i at most from S_i to r_i , ensuring that each source obtains its fair share of the CDR. In some embodiments, to achieve a rate of at least CDR, the WAN **1** and/or SPN **500** may at its discretion transmit traffic from a given active source or sources at a rate greater than r_i . In fact, the WAN **1** and/or SPN **500** may at its discretion transmit data from a source i at any rate between and including the fair share rate r_i and the full offered rate S_i .

If S_i is greater than T_i , a source may be considered by the WAN **1** and/or SPN **500** to be a non-conforming source. Conformance of a source may be calculated using a standard

leaky bucket algorithm with variable drain rate. Thus, the conforming depth of a bucket would be $DRS_i * W$. In other words, the maximum number of bits that will be sent to the network within a given time window of length W is equal to $DRS_i * W$. During a given time window of length W , the drain rate of the bucket is equal to T_i which is calculated during previous time windows. Thus, data packets inserted above the conforming bucket depth may be labeled as non-conforming. In other words, for a given time window, data packets in excess of the total $DRS_i * W$ number of bits may be labeled as non-conforming data packets. In such a situation, some or all of the source data packets equal to the difference between S_i and T_i may be labeled as non-conforming data packets, and some or all of the non-conforming data packets may be dropped.

This does not mean that data cannot be of a bursty or rate-variant nature. Although exemplary embodiments have been described as operating using average rates, real-time rates may vary within any given time window of length W . Thus, a certain amount of burstiness of data is allowable. This maximum burst size is the maximum number of bits that the WAN 1 and/or SPN 500 guarantees to transfer during a time window W .

In further embodiments of the invention, the WAN 1 and/or SPN 500 may provide forward congestion notification to a destination. For example, the WAN 1 and/or SPN 500 may provide a layer 2 binary indication that the CDR is being exceeded by using the frame relay forward explicit congestion notification (FECN) bit and/or a layer 3 message that indicates a non-conforming source and optionally contains rate information for that source (e.g. the actual transmitted rate T_i and/or the excess rate $S_i - T_i$). Furthermore, in some embodiments, multiple non-conforming sources might be listed, even within a single message. In these forward congestion notification embodiments, conformance may be measured at the network side of a destination. In some embodiments, a forward congestion notification may be provided to a given destination when the offering rate S_i of an active source offering to send traffic to the destination exceeds the actual network transmission rate T_i for the source.

Non-conforming packets that cannot be transmitted on the egress port of a source may be dropped with or without any indication to the source or destination. To measure conformance of a source, the amount of excess bandwidth available to the sources for transmission to the destination should be determined. To calculate the excess bandwidth, let W_j be the j^{th} time window. The excess bandwidth above the fair share bandwidth may be computed as

$$E = CDR - \sum_i \min(r_i, S_i) - MB,$$

wherein M is defined as the number of possible sources from which a destination may receive traffic, and wherein B is defined as a predetermined reference rate. The introduction of reference rate B effectively reserves network bandwidth for an inactive source, thus ensuring that a previously inactive source that becomes active can send at least some traffic through the network during time period W_j . Specifically, the WAN 1 and/or SPN 500 may ensure that each source T_i is guaranteed to be at least a minimum reference rate B . In this situation, a source is considered active during W_j if more than $B * W_j$ units of data (e.g., bits) are received during W_j . It is desirable to define B to be relatively small as compared with S_i so as to retain as much excess bandwidth as possible, yet

still large enough to ensure network availability to a non-active source (non-sending source with respect to a given destination) that may later become active with respect to a given destination. In some embodiments, B may be a predetermined rate. In further embodiments, B may vary with time, with the number of inactive sources, with the number of active sources, and/or with the total number of sources. In still further embodiments, B for a source may depend upon a priority classification assigned to the source. In still further embodiments, when a previously inactive source becomes active, the priority assigned to the source may depend upon the content of the data (e.g., data payload, DLCI, and/or address) offered to be sent. Thus, B may not be the same for each source.

Once the excess bandwidth is determined, the maximum conforming actual network transmission rates, T_i , may be calculated. To accomplish this, T_i for each source may first be set by default to $\min(r_i, S_i)$. Then the excess bandwidth, E , may be distributed among some or all of the sources that are actively transmitting to the given destination, thus adjusting or raising T_i for these sources. In some embodiments, the excess bandwidth may be uniformly distributed among some or all of the active sources. In further embodiments, the excess bandwidth may be distributed among these sources according to source priority, data priority, and/or DLCI.

In further embodiments, the WAN 1 and/or SPN 500 may provide backward congestion notification to a non-conforming source. Such notification may be in the form of a layer 2 and/or a layer 3 message indicating a destination(s) for which the non-conforming source is exceeding T_i and/or rate information for the non-conforming source (e.g. the actual transmitted rate T_i and/or the excess rate $S_i - T_i$). However, a layer 2 notification by itself may not be preferable, since a source receiving only a layer 2 notification may not be able to distinguish between destinations to which the source is conforming and those for which it is not conforming. In some embodiments, a backward congestion notification may be provided to a given active source when the offering rate S_i of the source exceeds the actual network transmission rate T_i for the source. In further embodiments, a user at a non-conforming source may be notified of congestion information, the assigned CDR, DRS_i , r_i , and/or T_i . In still further embodiments, it may be up to a user to decide how to act upon a congestion notification. In even further embodiments, a source may reduce its offering rate S_i in response to receiving a backward congestion notification.

In these backward congestion notification embodiments, conformance may be implemented at the network side of the source UNI. In such embodiments, feedback concerning the destination delivery rate may be required from the destination. The feedback may also contain information regarding the rate share of the active sources at the destination and/or the CDR divided by the aggregate rate.

While exemplary systems and methods embodying the present invention are shown by way of example, it will be understood, of course, that the invention is not limited to these embodiments. Modifications may be made by those skilled in the art, particularly in light of the foregoing teachings. For example, each of the elements of the aforementioned embodiments may be utilized alone or in combination with elements of the other embodiments. Additionally, although a meshed network is shown in the examples, the inventions defined by the appended claims is not necessarily so limited. Further, the IP switch may convert from any higher level IP like protocol to any fast-packet like protocol and is not necessarily limited to the ATM/IP example provided above. Furthermore, examples of steps that may be performed in the implementa-

tion of various aspects of the invention are described in conjunction with the example of a physical embodiment as illustrated in FIG. 5. However, steps in implementing the method of the invention are not limited thereto. Additionally, although the examples have been derived using the IP protocol for layer three, it will be apparent to those skilled in the art that any version of IP or IPX could be used as the layer three routeable protocol. Furthermore, it will be understood that while some examples of implementations are discussed above regarding IP and ATM protocols, the invention is not intended to be limited solely thereto, and other protocols that are compatible with aspects of the invention may be used as well.

We claim:

1. In a fast-packet network, a method comprising: managing, via one or more network elements, according to a committed delivery rate (CDR) at least one of a plurality of actual network transmission rates for at least one of a plurality of active sources, the committed delivery rate being associated with a destination, wherein the managing includes controlling a total delivery rate (R) to the destination according to the committed delivery rate (CDR) and a plurality of offering rates (S) of a first group of the plurality of active sources i, the active sources in the first group offering to send a plurality of data packets to the destination, such that:

$$R \geq CDR \quad \text{if } \sum_i S_i \geq CDR;$$
$$R = \sum_i S_i \quad \text{otherwise.}$$

2. The method of claim 1 wherein the managing further includes identifying at least one of the plurality of data packets as being nonconforming when a sum of the offering rates of the first group of active sources is greater than the committed delivery rate.

3. The method of claim 2 wherein the managing further includes dropping at least one of the identified data packets.

4. In a fast-packet network, a method comprising: managing, via one or more network elements, according to a committed delivery rate (CDR) at least one of a plurality of actual network transmission rates for at least one of a plurality of active sources, the committed delivery rate being associated with a destination; and assigning a destination rate share (DRS) to each of a first group of the active sources, the first group of active sources offering to send a plurality of data packets to the destination, wherein said managing includes managing according to the destination rate shares of the first group of active sources an actual network transmission rate (T) for at least one of the active sources in the first group of active sources, wherein the managing further includes

determining a fair share rate (r) for at least one of the active sources i in the first group of active sources according to the destination rate share (DRS) of the at least one active source and the committed delivery rate (CDR), such that:

$$r_i = \frac{DRS_i}{\sum_i DRS_i} CDR.$$

5. The method of claim 4 wherein the managing further includes adjusting the actual network transmission rate (T) for at least one of the active sources i in the first group of active sources according to an offering rate (S) of the at least one active source, the fair share rate (r) of the at least one active source, and the committed delivery rate (CDR), such that:

$$\text{when } \sum_i S_i \geq CDR,$$
$$T_i \geq \min(r_i, S_i).$$

6. The method of claim 5 wherein the managing further includes identifying at least one of the data packets as being nonconforming when a sum of the offering rates of the first group of active sources is greater than the committed delivery rate.

7. The method of claim 6 wherein the managing further includes dropping at least one of the identified data packets.

8. In a fast-packet network, a method comprising: managing, via one or more network elements, according to a committed delivery rate (CDR) at least one of a plurality of actual network transmission rates for at least one of a plurality of active sources, the committed delivery rate being associated with a destination, wherein the managing includes notifying at least one of the active sources of a network congestion by providing a backward congestion notification to the at least one active source when an offering rate of the at least one active source exceeds the actual network transmission rate for the at least one active source, wherein the providing a backward congestion notification to the at least one active source includes providing a Layer 3 backward congestion notification to the at least one active source, wherein the providing a backward congestion notification to the at least one active source further includes providing information representing an identity of the destination, the at least one active source offering to send a plurality of data packets to the destination.

* * * * *